MULTITAPER SPECTRAL ESTIMATION AND TIME-DOMAIN CROSS-CORRELATION IN FMRI DATA ANALYSIS: ACTUAL AND SIMULATED DATA

By

WAYNE M. KING

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

1999

For Henry: nothing I say here or in the future could ever be more important than our time together, and for Vernon, whose gentle push started the journey.

ACKNOWLEDGMENTS

Suffice it to say that this endeavor would not have been possible without the help and inspiration of many people. I would like to thank Partha P. Mitra, David J. Thomson, Donald B. Percival, and Andrew T. Walden for their very patient and kind help in facilitating my understanding of multitaper spectral analysis. Partha introduced me to multitaper spectral analysis two years ago and has been incredibly helpful and generous both with his work and time. If it were not for him, it is doubtful that I would have found a dissertation topic I enjoyed so thoroughly. David J. Thomson and Donald B. Percival also deserve special recognition for their exceedingly gracious help. Both have gone to great lengths to make themselves available by phone and e-mail with references, answers, and encouragement in spite of having never met me and their own demanding schedules. They have provided me with a shining example that I promise to never forget.

Closer to home I would like to thank Christiana M. Leonard for her support, both financial and academic, over these last several years. Dr. Leonard cannot possibly know how many students, including myself, she has encouraged with her scientific rigor and integrity and touched with her kindness. It was nice to be part of a lab that felt like a family. I would like to thank Stephen J. Blackband for granting me unselfish use of his lab and computers even though I was doing fMRI and not "real science". My understanding of the amazing world of nuclear magnetic resonance, however inadequate and incomplete, owes much to him. My own developing passion for mathematics has been nurtured and

encouraged by Professor Murali Rao's love for the subject and his complete humility in the face of its depth. Professor, I promise to always treat it with respect. Linda J. Lombardino has a special gift for looking beyond academics and recognizing the human being inside of the student. There are many of us who have benefited from her advocacy, compassion, and sense of humor. Kenneth Gerhardt has never ceased to provide me with an example of what can be achieved with scientific rigor, passion, and perseverance.

Finally, I would like to thank the two people, who have been there for me every step of this journey. First, Cindy, who never failed to provide me whatever financial and emotional support I needed, in spite of receiving precious little in return. And Carl C. Crandell, who has juggled the difficult roles of chairman and friend, while somehow managing to perform each of them equally well.

TABLE OF CONTENTS

page

ACKNOWLEDGEMENTS....iii

LIST OF TABLES....vii

LIST OF FIGURES....viii

ABSTRACT....x

1 MRI BASIC PRINCIPLES....1

1.1 NMR-A Quantum Treatment....3
1.2 From Waves to Particles and Back to Waves....6
1.3 Spin....7
1.4 Transverse Magnetization....12
1.5 Spin-lattice and Spin-Spin relaxation....13
1.6 T2* Effects....15
1.7 Magnetic Susceptibility....16
1.8 Gradients....17
1.8.1 Slice Selection....18
1.8.2 Frequency-Encoding Gradients....21
1.8.3 Phase-Encoding Gradients....22
1.9 Pulse Sequence Design (PSD)....24
1.10 Image Weighting....29
1.11 Fast MRI....31
1.11.1 Echo-Planar Imaging....32
1.12 Diffusion and Perfusion in MRI....33
1.12.1 The BOLD Contrast....37
1.12.2 Arterial Tagging Methods....39
1.13 Brain vs. Vein Debate....40

2 FUNCTIONAL IMAGING-PHYSIOLOGICAL BASIS....42

2.1 Brain Activation....42
2.2 Cellular Components of the Mammalian Brain....44
2.3 The Synapse....45
2.4 Neurotransmitters....46
2.5 Energy Metabolism in the Neuron....47

2.6 Anatomy of Cerebral Circulation....49
2.6.1 Microvascular Organization....50
2.6.2 Regional Cerebral Blood Flow (rCBF)....50
2.6.2.1 Nitric oxide....52
2.6.2.2 Potassium....52
2.6.2.3 Adenosine....53

3 TIME SERIES IN FMRI....56

3.1 Stochastic Processes....58
3.2 Time/Frequency Analysis....61
3.2.1 The Cross-Correlation Method....68
3.2.2 Frequency-Domain Analysis....71
3.2.3 The Autocovariance Sequence....72
3.2.4 Non-Parametric Spectral Estimation....75
3.2.5 The Periodogram....77

4 MULTITAPER SPECTRAL ESTIMATION....83
4.1 Slepian Sequences....83
4.2 The Multitaper Spectral Estimator....86
4.3 Line Detection in the Multitaper Method....89

5 ANALYSIS OF SIMULATED AND ACTUAL FMRI DATA....94
5.1 Application to Simulated Data....94
5.2 Application to Actual fMRI Data....101
5.2.1 Visual Experiment....103
5.2.2 Motor Experiment....118
5.3 Discussion....123
5.4 Future Directions....130

REFERENCES....132

BIOGRAPHICAL SKETCH....136

LIST OF TABLES

Table page

1 Common windowing functions. 80
2 Significant correlations by sample size. 100
3 Correct detections by sample size. 100

LIST OF FIGURES

Figure		page
1	Flip angle illustration	25
2	Pulse Sequence Diagram	33
3	Neuronal basis of the BOLD contrast	54
4	550-point image xcorr	105
5	550-point image MTM	106
6	150-point image xcorr	108
7	150-point image MTM	109
8	250-point image xcorr	110
9	250-point image MTM	111
10	350-point image xcorr	112
11	350-point image MTM	113
12	450-point image xcorr	114
13	450-point image MTM	115
14	500-point image xcorr	116
15	500-point image MTM	117
16	XCR activation map-TR=3	119
17	MTM activation-TR=3 seconds	120
18	XCR activation right-finger tapping	121

19 MTM activation right-finger tapping.....122

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Doctor of Philosophy

MULTITAPER SPECTRAL ESTIMATION AND TIME-DOMAIN CROSS-CORRELATION IN FMRI DATA ANALYSIS: ACTUAL AND SIMULATED DATA

By

Wayne M. King

August 1999

Chairman: Carl C. Crandell, Ph.D.
Major Department: Department of Communication Sciences & Disorders

Functional Magnetic Resonance Imaging (fMRI) is an important emerging tool in the study of sensory and information processing in the human brain. However, at present there is no general agreement on the best method for the analysis of fMRI data. A major reason for this lack of agreement is the paucity of methodological studies that have undertaken systematic comparisons of competing data analysis methods on actual and simulated fMRI data. This lack of comparative methodological studies is unfortunate because the developmental of fMRI as a reliable experimental and clinical tool depends crucially upon theoretically sound and proven data-analysis methods.

With these considerations in mind, this project compares two competing data

analysis methods on actual and simulated fMRI data. The actual fMRI data was obtained in visual and motor experiments conducted at 3-Tesla using single-shot Echo Planar Imaging (EPI). The simulated data was generated to closely approximate signal-to-noise (SNR) ratios and noise variances observed in actual fMRI data. The two competing data-analysis methods used here are multitaper spectral estimation, a frequency-domain approach, and time-domain cross-correlation. Both methods are univariate procedures, which assume an experimental paradigm with a periodic forcing frequency.

The simulated data reported here suggest that the multitaper spectral estimator is more robust, both in terms of correct detections and false positives than the cross-correlation estimator. This results are corroborated by the actual data, which demonstrate that the multitaper method provides activation maps, which are more localized to the expected areas of neuronal activation.

CHAPTER 1
MRI-BASIC PRINCIPLES

Nuclear magnetic resonance (NMR), as the name suggests, refers to an enhanced absorption of energy by nucleons when they are irradiated with the proper frequency of RF energy. In turn, this absorbed energy is then emitted by the particles and captured as the fundamental building block of all modern NMR signal-processing applications, including NMR spectroscopy and magnetic resonance imaging (MRI). A key advantage of the resonance effect is that a sub-population of nucleons can be selected while effectively ignoring the rest, even though the contribution of the sub-population may be relatively weak (Slichter, 1978). For example, the emissions of nuclei would be impossible to separate from the much larger effects of electrons, if we were not able to utilize the principles of resonance. In this chapter the basic principles of NMR and MRI will be examined in some detail in order to more fully understand how these methods are used to investigate physiologic changes in the brain related to cortical activation. The approach will be chiefly classical mechanical, but it should be noted that the use of classical physics in the atomic and sub-atomic world results in paradoxes that can only be rectified with a quantum or statistical mechanical approach. Classical approaches to NMR are necessarily macroscopic, although there will be occasion to make reference to the microscopic, and hence, quantum principles at work. While simplistic, the classical approach is not without merit, indeed Bohr's correspondence principle mandates that quantum formulations agree

with the classical as the quantum numbers become very large. Finally, this treatment is unavoidably incomplete and sacrifices many of the details in hopes of acquiring the essential ingredients to understanding functional MRI (fMRI).

In summary, the follow steps are involved in the production of a signal utilizing nuclear magnetic resonance. Steps 1 & 2 are necessary for both imaging and non-imaging applications, e.g. spectroscopy, while steps 3 & 4 apply only to magnetic resonance imaging (MRI).

1. Randomly-oriented nuclei possessing *intrinsic spin angular momentum* are aligned by the application of an external magnetic field such that vectorially the individual magnetic moments sum to create a net magnetization vector, μ_o along the axis of the main external field, B_o.
2. This longitudinal magnetization must be "tipped" away from the main field in order for us to be able to detect it. This is a result of the fact that B_o is typically several orders of magnitude greater than μ_o. This "tipping" is achieved by resonant effects produced in the nuclei through the application of carefully-designed RF pulses. The actual frequency of RF energy required to induce the resonance effect is proportional to the external magnetic field and can be determined by the Larmor equation.
3. In order to obtain spatial information, the positions of the tipped nuclei are encoded by applying mutually-orthogonal linear gradients. The gradients produce spatially-dependent magnetic fields and, as a result, resonant frequencies which are also spatially-dependent.
4. The signals produced by the emission of energy by the nuclei are read in a

user-determined way to produce images with varying degrees and types of contrast.

Each of these points will now be examined in turn.

1.1 NMR-a Quantum Treatment

The concept of energy transfer, absorption and emission is critical to many everyday physical phenomena of both the visible and invisible universe. This is true whether one considers pushing a child on a swing, or the complicated interactions of molecules and RF energy in NMR. For purposes of illustration, consider the example of a pendulum. When we set a pendulum in motion, it begins to oscillate at its natural frequency ν, which depends on its length and is inversely proportional to the square root of its mass, $1/\sqrt{m}$. According to classical Newtonian mechanics, the amplitude of the pendulum's motion can increase without bound in a continuous manner as more and more force is applied. In other words, energy is seen to be a continuously distributed quantity, not restricted *a priori* to any particular values. However, early experiments in atomic particles contradicted these classical predictions. On the contrary, in the microscopic world it appeared that energy came in discrete increments; in other words, it was *quantized*. Theory and empirical evidence have converged to show that energy in our pendulum example can only assume values that are multiples of $h\nu$, where h is Planck's constant, 6.6×10^{-34} Joules-seconds. The reason that energy appears to be continuously-distributed on a macroscopic level is apparent when the magnitude of Planck's constant and the natural frequencies of large bodies are considered. The product of Planck's constant and the low natural frequencies of macroscopic bodies makes the interval between successive energy states extremely small, so small that one can consider them as continuous without any

appreciable loss of accuracy. However, in the atomic and sub-atomic world, where the masses encountered are infinitesimal, natural frequencies in the mega- and gigahertz range are readily encountered. With these bodies successive energy states are well-separated. Therefore, to account for mechanical interactions on the atomic level, energy must be treated as a discrete quantity. This is the fundamental point of departure for a quantum mechanical treatment of nature.

On a molecular level, we can distinguish three general types of motion important in the absorption and transmission of energy: translational, vibrational, and rotational (Atkins, 1986). Translational motion refers to the random motion molecules undergo in a container, or a sample. These translational movements impart kinetic energy to the sample, or container. Further, molecules can rotate and vibrate. In accordance with quantum mechanical predictions, allowable energy levels for each of these motions are discrete, and are represented by quantum numbers.

Translational, rotational, and vibrational motion in classical mechanics can be described in terms of a few basic equations. One of these expresses the total energy of a particle in terms of its kinetic and potential energy:

$$E=1/2mv^2 + V$$

where v is the velocity of the particle, m is its mass, and V is the potential energy. This equation may be written in terms of momentum if one considers that *linear momentum*, p is equal to the product of mass and velocity, $p=mv$. This yields the equation:

$$E=p^2/2m + V$$

It turns out that this form of the equation is much more useful in the solution of real problems. For example, since $p = m\frac{dx}{dt}$, we have a differential equation for x (the body's position) as a function of t, time. Accordingly, its solution will allow one to predict the trajectory of a particle if its position and momentum are known. The second fundamental equation of motion is Newton's second law:

$F=dp/dt$, or equivalently $F=ma$ where a=acceleration

This equation is easily derived from $p=mv$ if we properly consider momentum, p(t), and velocity (dx/dt) to be functions of time. Then we immediately see:

$$p(t)=mv(t)$$
$$p(t)=m(dx/dt)$$
$$dp/dt=m(dx^2/dt^2)$$

For angular motion, the expression becomes

$$J=I\omega$$

Where J is angular momentum, I is the moment of inertia and ω refers to angular velocity or frequency. Finally in the case of harmonic oscillation we have $F= -kx$, where a particle experiences a restoring force proportional to its displacement. Here k here is the *spring constant* which depends on the material being displaced. Recalling Newton's 2nd law, we have:

$$F=ma$$
$$ma= -kx$$
$$m(dx^2/dt^2)= -kx$$

The general solution to this second-order differential equation is given by

$$X(t)=A \sin \omega t + B \cos \omega t$$

In this case the position of the particle is seen to vary in a sinusoidal manner as a sum of two trigonometric functions with a frequency of $\omega/2\pi$. In summary, the equations of motion in classical physics: 1.) predict a precise trajectory and position if the necessary initial and or boundary conditions are known, and 2.) they allow the translational, rotational, and vibrational modes of motion to be excited to any degree, i.e. the amplitude in each case is continuous. As previously mentioned, this latter point is a major failing of classical physics.

1.2 From Waves to Particles and Back to Waves

At the beginning of the 20th century, the prevailing thought was that light was a wave. However, experiments in the early 1900s demonstrated that light also had particle-like attributes. The photoelectric effect and Compton scattering are two prime examples of the particle nature of light. A closer examination of this apparent paradox led to one of the mainstays of modern physics: when events on an atomic scale are examined, the properties of particle vs. wave become obscured, i.e. particles show wave-like properties and waves act like particles. In 1924, de Broglie helped to eliminate this paradox by suggesting that any particle moving with momentum, p, should have a wavelength given by:

$$\lambda = h/p$$

From this key discovery, physicists began to abandon the classical concept of a localized particle and adopted the view that a particle's position is distributed like the amplitude of a wave, i.e. with each particle one could associate a wavefunction. In 1926, Erwin Schroedinger proposed an equation that yields the wavefunction for any system when solved (Atkins, 1986). Furthermore, wavefunctions must be elements of a normed vector space and their square of the inner product with themselves must satisfy the properties of a

probability density function (pdf) (unit norm). Now, in contrast to the equations of motion in classical theory discussed above, quantum theory predicted that a particle's position could only be stated probabilistically and energy states were restricted to a small number of discrete values in agreement with their Schroedinger equation. Also, in quantum theory, position and momentum are Fourier transforms of one another, implying that when position is completely determined momentum is completely undetermined and vice versa.[1] This is known as *Heisenberg's uncertainty principle.*

1.3 Spin

We have already discussed that a particle may exhibit rotational motion, or angular momentum. Additionally, we have seen that the classical equation for angular momentum fails because it does not predict the discrete nature observed on the atomic level. However, what surprised even the early quantum theorists was that particles can exhibit more than one kind of rotational motion. An electron for example exhibits an *orbital angular momentum* as it travels around its nucleus. Mathematically, this orbital angular momentum can be modeled as a point orbiting the surface of a sphere. Furthermore, the solution of the wavefunction predicts that the orbital quantum number can only assume integer values (Atkins, 1986). However, a famous experiment performed in 1921 by Otto Stern and Walther Gerlach seemed to contradict this postulate. Stern and Gerlach projected silver atoms through a non-homogenous magnetic field and observed the bands produced by collisions of the atoms on a screen outside the magnetic field. When the particles were projected through sufficiently slowly, the silver atoms segregated into two

[1] This is also true of the time and frequency domains of a signal' representation and this point will be illustrated later in terms of an RF pulse.

distinct bands. In some sense this result was comforting since it supported the quantum notion of discrete orientations and contradicted the classical predictions. Yet, in another sense it was puzzling since two bands could only be obtained if the orbital angular momentum were allowed to assume half-integer values, namely ½. This was in contradiction to the conclusion that orbital angular momentum must be an integer. In the end, this contradiction was resolved by discovering that what they were observing was not due to the atom's orbital angular momentum, but rather due to its own intrinsic angular momentum, or *spin*. The fact that intrinsic spin numbers can assume half-integer values is a result of not having to satisfy the same boundary value conditions as a particle moving over the surface of a sphere (orbital angular momentum). For an electron, this *intrinsic angular momentum*, or *spin* can assume exactly one of two values, +1/2, or spin up, also known as an α-electron, or the spin down, -1/2 state. The latter is sometimes referred to as a β-electron. The notion of spin is critical in NMR, because it is this property which makes particles visible to magnetic resonance methods. In addition to electrons, some nuclear particles also possess spin. Since the nucleus is an aggregate of more elementary particles, one might expect the picture of spin to be much more complicated on a nuclear level. This is indeed the case. Yet, we can make some very simple generalizations: 1.) if the nucleus has an even number of protons (atomic number) and an even atomic weight, it will have no spin. Such a nucleus will not interact with an applied magnetic field and cannot be studied using NMR. 2.) If a nucleus has an even atomic weight and odd atomic number it will have an integer value for its spin. 3.) If a nucleus has odd atomic weight, it will have a half integer value for its spin. Nuclei with spins possess magnetic moments. Further, a nucleus with spin S may take $2S+1$ orientations in a magnetic field. These

orientations are given by S, S-1,… -S. Thus, as previously mentioned, an electron with spin ½ can assume one of two possible orientations, either +1/2 or –1/2. The energy difference between these two states, ΔE, is proportional to the applied magnetic field. If we irradiate a sample with energy equal to ΔE ($h\nu = \Delta E = |E_\beta - E_\alpha|$) then a strong coupling will result between the spins and the applied radiation. This coupling is analogous to the effect seen in macroscopic bodies if force is applied at their natural frequencies. On an atomic level, the effect of this coupling is to induce transitions in some of the spins from $\alpha \rightarrow \beta$. It is precisely this effect that forms the basis of all NMR methods.

There are many nuclei, which are available for the type of interaction just described. However, H^1 is the nucleus of choice for many imaging applications, including fMRI, for a number of reasons. The most important of these is the high natural abundance of hydrogen in the human body and its large magnetic moment. The former is due to the large presence of water and fat , while the latter relies on the fact that the magnetic moment of a nucleus is inversely proportional to its mass. This can be seen in the formula for the magnetic moment of a nucleus called the *Bohr nuclear magnetom*.

$$\mu = e\hbar / 2m$$

In the above equation, μ is the total magnetic moment, h-bar is Planck's constant divided by 2π and m is the mass of the proton. Since H^1 with one proton in its nucleus has the smallest mass of any element on the periodic table, it can be expected to have the largest magnetic moment. Additionally, hydrogen has an odd atomic weight which means it possesses a half-integer spin number (actually it possesses spin 1/2 which is desirable for reasons beyond the scope of our discussion).

Armed with notions of spin, the discrete nature of energy, and resonance, we are ready to discuss how an observable signal can be produced using NMR. Assume that we have placed a large sample of hydrogen protons in a magnetic field. This large sample could be part of a living organism, or simply a bottle of water. We know that H^1 possesses spin and will accordingly interact with an external magnetic field. Adopting a classical point of view, we can regard each of these tiny hydrogen protons as a bar magnet with a north and south pole. In most materials, these bar magnets are randomly aligned so that when one considers the vector sum, the net magnetic moment is zero. In order to create the necessary initial conditions for nuclear magnetic resonance, we must somehow induce a net magnetic moment in the material. This is done by the application of an external magnetic field. When nuclei possessing spin are exposed to an external magnetic field, they align themselves either parallel or anti-parallel to the applied field, with parallel alignment being a lower energy state. In quantum theory this is referred to as the *Zeeman effect*, or the *Zeeman interaction* (Merzbacher, 1961). According to results from classical mechanics, a magnetic moment in a magnetic field will experience torque as it rotates around the magnetic field vector, B_o. This torque will guarantee that the magnetic moments of the individual spins "precess" around the main magnetic field much like a top will precess under the influence of gravity. This precession effect could actually be seen with an ordinary bar magnet in a magnetic field if the bar magnet could somehow be mounted on a frictionless system. If we examine in detail the effect of the torque produced by the interaction of a static magnetic field on a body with angular momentum, we can derive the important *Lamor* relationship, which states that a magnetic moment will precess at a frequency proportional to magnetic field strength given by $\omega_o = \gamma B_o$, with γ a

proportionality constant called the *gyromagnetic ratio.* With knowledge of the natural precessional frequency of a nucleus at a given magnetic field strength, B_o, we are in a position to induce resonance effects by introducing energy near the precessional frequency (Slichter, 1978). For the hydrogen proton the frequency for maximal energy absorption turns out to be approximately 42.58 mHz/Tesla(T), with Tesla being the derived SI unit for magnetic field strength. Therefore in a 1.5T field, presently the most common choice for commercial whole-body systems, the resonance frequency of H^1 is 63 mHz (approx.).

The introduction of RF energy into a spin system at the Lamor frequency results in the most efficient absorption of energy, or resonance. This energy absorption has two important side effects, which allow the spin system to be "visualized". First, energy absorbed by the system is released once the stimulation is stopped. This process involves the excited "tipped" spin system returning to its equilibrium position, aligned with the main magnetic field vector. Note that the spin system could be tipped trough 2π radians by a sufficiently powerful application of RF energy and consequently be brought "back" into equilibrium. However, this is not done in practice for reasons that will be made clear shortly. In NMR terms, RF energy at the Lamor frequency results in a highly-efficient coupling with the spin system and can induce transitions between spin states. Once the application of RF energy is stopped, the spin system will begin to release the absorbed energy and return to its pre-pulse state. Secondly, the absorption of RF energy by the spin system produces a short-term phase-locking between individual spins. As will be seen in the next section, this phase-locking is critical for the detection of the relatively weak magnetic moment of the spin system.

1.4 Transverse Magnetization

Having induced resonance effects via the Lamor relationship, we are in a position to detect our spin system as it releases the absorbed energy. However, in order to accomplish this, the spin system must somehow first be "tipped" away from the main magnetic field. Yet, when we look at the effect of resonance, it appears that we have only altered the longitudinal magnetization of the system by inducing spins to change their energy states. There does not appear to be any transverse magnetization. This apparent paradox is the result of attempting to explain the behavior of a group of protons by the action of a single member. We have already said that a properly-tuned RF pulse will cause the number of protons in the spin up and spin down states to change. Another effect that the RF pulse has is to cause the individual spins to precess in phase in the same way that an axle on a car locks the positions of two wheels (Wehrli, 1988). Since the magnetic moments are precessing around the main magnetic field in a wobbly fashion, when they are forced to precess in phase, their "wobbling" creates a net transverse component to the magnetization. The angle (or amount) of the transverse magnetization is dictated by

$$\Theta = \gamma B_1 t$$

where B_1 is the amplitude of the RF pulse and "t" is the duration. Θ here is referred to as the *flip angle*. As can be seen from the above relationship, we can alter the flip angle by manipulating the amplitude or duration of the RF pulse.

Once we have created a coherent transverse component to our magnetization vector μ, we can "see" the spin system with the help of Faraday's law of electromagnetic induction. Recall that a moving electric charge creates a magnetic field and a fluctuating magnetic field induces electric charge to flow. By creating a coherently precessing transverse

magentization, we have created the latter situation. Now if we place a properly tuned wire in the vicinity of this fluctuating magnetic field, by classical electromagnetic theory, we should be able to induce a current in the wire. This is in fact exactly what happens and this induced current takes the form of an exponentially-decaying sinusoid called the *free induction decay* (FID). The reason for the sinusoidal nature of the signal should be clear from the fluctuating nature of the magnetic fields that induced it. Why it should decay exponentially is the subject of our next discussion.

1.5 Spin-Lattice and Spin-Spin Relaxation

Relaxation can be broadly defined as the evolution toward statistical equilibrium of macroscopic systems (Fullerton, 1988). Once the introduction of RF energy into the spin system is stopped, the coherent precessing that created the transverse magnetization will begin to decay. It was Felix Bloch who first proposed that the return of this magnetization to equilibrium was exponential and followed two time constants, which he called T1 and T2. He designated the longitudinal relaxation by T1 and the transverse relaxation by T2. These processes are also known as spin-lattice and spin-spin relaxation respectively. Spin-lattice relaxation refers to the process by which the excited spins exchange energy with the atomic lattice in which they are located until thermal equilibrium is reached. In the case of the energy levels encountered in NMR we can regard the lattice as an infinite medium whose own thermal equilibrium is not affected by exchange of energy with the spin system (Fullerton, 1988). Spin-spin relaxation is the process by which spins interact with one another to cause a dephasing, or loss of coherence in the transverse magnetizaton.

To fully understand these processes requires a detailed journey through quantum and statistical mechanics. Specifically, one must consider the ensemble average of the mean

square departure produced by each individual spin's precession. For our purposes it will suffice to examine them in a more qualitative way. All molecules are in a constant state of motion. As they move, they collide or come into close proximity with other molecules. Since these molecules consist of moving charged particles, they have tiny associated electric and magnetic fields. These field interactions between molecules cause fluctuations in the local magnetic fields experienced by each. These interactions are referred to as *coupling*. A macroscopic example of this coupling can be observed if two magnetic compasses are placed adjacent to one another. In close proximity, the compasses react not only to the magnetic field of the earth, but also to the magnetic field of the other compass. If the needle of one compass is moved, the needle of the other compass will move in response. The needles of the two compasses are coupled magnetically.

On a microscopic level, any source of a magnetic field can be a possible source of coupling. However, in the case of hydrogen protons, the most important type of coupling is the interaction between neighboring dipoles, or the so-called *dipole-dipole coupling*. Specifically, the time that the dipoles and their associated magnetic fields are in contact with each other, known as the *correlation time*, τ_c determines how substantial the dipole-dipole interaction is. In other words, how quickly relaxation takes place. Basically, the longer spins are in contact with one another, the more effectively they are able to transfer energy as well as affect the phase coherence of the system. In this way, both spin-lattice and spin-spin relaxation are promoted (Fullerton, 1988). The amount of time they spend in contact crucially depends on how rapidly their magnetic fields are fluctuating, or alternatively, how fast they are spinning. Furthermore, how fast a molecule spins depends on its chemical environment. For example, the hydrogen protons in free water are able to

move more easily than protons in fat (hydrocarbon chains) or in other macromolecules, such as proteins. Therefore, in pure water, the fluctuations are quite rapid and hence the time that the molecules are in contact with one another becomes very small. As a result, there is little promotion of relaxation of either kind.

In data collection we are always limited by the relaxation properties of our tissue, since they cause an exponential decay of our signal. Chiefly, we are limited by the T2 of the tissue, since $T2 \leq T1$. The reason for this is intuitively clear. The production of a FID depends upon the coherent precesssion of spins, thereby creating a transverse magnetization vector. Once this transverse magentization has decayed, we will no longer be able to induce a current in our receiver coil. The loss of this phase coherence typically occurs (for most materials) faster than a complete return to longitudinal equilibrium. In some materials, such as solids, the T2 is so short that imaging is extremely difficult.

1.6 T2* Effects

Modern superconducting magnets have remarkably uniform magnetic fields. However, whenever a volume of tissue (or any substance) is placed in these magnets, some inevitable differences in the magnetic field develop within that tissue. The differing magnetic susceptibilities (see below) that inevitably exist within the sample are largely responsible for these inhomogeneities. The effect in terms of NMR is that these static (not due to molecular motion) inhomogeneities contribute to transverse (spin-spin) relaxation. As a result, spin-spin relaxtion occurs due to dipole-dipole coupling plus the added component of static field inhomogeneities. This combined effect is referred to as T2* (T2-star).

1.7 Magnetic Susceptibility

The general term for the magnetizability of a material is its *magnetic susceptibility,* χ. (Purcell, 1965). χ is a dimensionless quantity usually expressed in parts per million. In NMR magnetic susceptibility effects typically show up as a local loss of signal that results from an accelerated dephasing of the spin system. This fact is troublesome in some instances, but it is also exploited in the use of exogenous and endogenous contrast agents. In fact, the magnetic susceptibility differences between oxygenated and deoxygenated hemoglobin is the basis of the most popular current method of fMRI, the **B**lood **O**xygen **L**evel **D**ependent contrast. In view of this, it will be instructive to gain a deeper understanding of what magnetic susceptibility actually is.

Of all the atomic constituents that interact with a magnetic field, the nuclear interaction is one of the weakest. Much larger interactions for example occur with electrons, which have magnetic moments several orders of magnitude greater than protons. Generally speaking, materials can be classified in one of three ways, depending on how they react to a magnetic field: diamagnetic, paramagnetic, and ferromagnetic. These classifications are chiefly based on the presence or absence of unpaired electrons. *Diamagnetic* materials are those which tend to resist a magnetic field. Diamagnetic materials have no unpaired electrons. On the other hand, *paramagnetic* materials have unpaired orbital electrons. These electrons tend to line up with field producing additive fields of their own. This effectively increases the strength of the magnetic field at a given point. The rare-earth element gadolinium (Gd) with seven unpaired electrons is a strongly paramagnetic substance often used as a contrast agent in MRI. Deoxyhemoglobin has four unpaired electrons and is paramagnetic, while its oxygenated counterpart, oxyhemoglobin is

diamagnetic. Finally, *ferromagnetic* substances are those that remain permanently magnetized even after they are removed from an applied field. By virtue of producing additive fields, paramagnetic substances can accelerate the rate of T2* decay. As alluded to above, this accelerated dephasing can be detected as a focal loss of signal in an image. This fact can be used to determine which regions of an image contain signficant proportions of either an exogenous or endogenous paramagnetic substance. Ferromagnetic materials can also produce this loss of signal but their usage is not as common.

1.8 Gradients

In order to create an image using magnetic resonance we need to obtain spatial information from the spin system. In other words, spins located in different parts of a volume of tissue must somehow be encoded differently. This is accomplished by exploiting the fact that the resonant frequency of a proton is dependent on the strength of the magnetic field that it experiences. This leads to the conclusion that if the magnetic field is made spatially-dependent in some manner, it would be possible to distinguish protons which inhabit different locations in the sample. In the language of NMR, these spatially dependent magnetic fields are called *gradients*.

The spatial encoding of spins most common in modern MRI is called *spin-warp imaging*. Spin-warp imaging utilizes three orthogonal gradients to accomplish spin localization in two spatial dimensions: 1.) the slice selection gradient (Gz), 2.) the frequency or readout gradient (Gx), and 3.) the phase encoding gradient (Gy). In most descriptions of the gradient system, arbitrary orientations are used for purpose of illustration. It should be noted, however, that the actual physical layout of these gradients

in any given NMR is typically much more complicated than is described in introductory texts or will be described here.

1.8.1 Slice Selection

In applications of MRI to biological tissues, we are typically not concerned with obtaining an image of the entire object. We may, for instance, be interested in only a small section of the human brain. This is often the case in functional imaging, where one may only be interested in signal changes in a particular sensory cortex. In order to avoid exciting all spins in the system, we have to somehow change the resonant frequency of the region of interest and then selectively shape our RF pulses to induce resonance only in that region. We first see how this shaping is accomplished with an RF pulse.

Let us consider an RF pulse which approximates a sinc function- sin x/x. An examination of the sinc function will reveal several things immediately

1.) The sinc function consists of a sine wave, whose amplitude is modulated by the function 1/x. As $x \Rightarrow \infty$, the envelope of the sinc function goes to zero.

2.) If we evaluate the sinc function as x approaches zero. We find using L'Hopital's rule that

$$\lim_{x \Rightarrow 0} \sin x/x = 1$$

3.) Finally since sin(x) is an odd function, sin(-x)= -sin (x). Accordingly, when sin(x) is negative, so is 1/x. Therefore sinc(-x)=sinc(x) and sinc(x) is an even function. We now see that sinc(x) has its greatest amplitude at x=0 and is symmetrical with respect to the y-axis.

Now let us examine the relationship between the sinc function and its Fourier transform cousin the rectangular pulse. Since the information in a function and its Fourier transform are equivalent, although in different domains, we will work from the rectangular pulse to the sinc function. This operation is accomplished by use of the inverse Fourier transform given by

$$f(t)=\int_{-\infty}^{\infty} F(\omega)e^{i\omega t}\,d\omega$$

Let us consider a rectangular pulse from –a to a with an area of 2a. The inverse fourier transform is then given by

$$f(t)=\int_{-a}^{a} F(\omega)e^{i\omega t}d\omega$$

$$f(t)=\int_{-a}^{a} e^{i\omega t}\,d\omega$$

$$1/it[\, e^{i\omega a}-e^{-i\omega a}]$$

Now recalling that $\sin\theta=e^{i\theta}-e^{-i\theta}/2i$. We can rewrite the above as

$$2a \sin at/at \text{ or } 2a \operatorname{sinc}(at)$$

Since sin (at)=0 when $at=n\pi$, where n is an integer then when $t=n\pi/a$ the function will be zero, with the first zero crossing coming at π/a. If "a" represents the duration of the sinc function, then the shorter the duration the greater the value of the first zero crossing. In a Fourier transform pair such as sinc function⇔rectangular gate pulse- the relationship is unaffected by which one is assigned to the time vs. frequency domain. A sinc-shaped pulse in time will generate a rectangular frequency profile. Similarly a rectangular pulse in time would generate a sinc-shaped frequency profile. In NMR, one is often interested in

obtaining a rectangular frequency profile in order to excite a particular portion of the volume of interest. This rectangular profile can be obtained by using an RF pulse which is sinc-shaped in time. Furthermore, if we shorten the duration of the sinc pulse, its frequency domain representation will be become broader. In this way we can use the FT relationship to tailor our RF pulses to excite a selective band of frequencies. Given that we can use the FT relationship to control the frequency content of our RF pulse, we must also be able to control the Larmor frequency of our sample. This is accomplished by the use of a slice-selection gradient.

As we have seen gradients are spatially dependent magnetic fields that are applied on top of the main static field. By applying a gradient before we excite our sample and by controlling the frequency content of our exciting RF pulse we can selectively excite only a portion of the spin system. As we have seen, the RF pulse is not comprised of a single frequency by rather by a band of frequencies, in other words, RF pulses can be constructed to be of a certain *bandwidth*. We can also increase or decrease the strength of the slice-selection gradient. If we increase the strength (amplitude) of a linear gradient, we increase its slope and thereby create a greater range of resonant frequencies across the profile of the gradient. An RF pulse with a given bandwidth will accordingly excite a narrower portion of the spin system when the slice-selection gradient amplitude is increased. In effect we will receive information from a narrower portion of the volume of interest. In terms of imaging this will result in a narrower slice. Thus by altering the amplitude of the slice selection gradient we can excite a volume of tissue of greater or lesser thickness. Note that we could also accomplish this by fixing the slice-selection

gradient and changing the bandwidth of the RF pulse, but most current commercial systems utilize the former method.

1.8.2 Frequency-Encoding Gradients

Once we have managed to selectively excite a particular slice of tissue, we need to be able to determine the positions of spins within that volume. This is accomplished by the utilization of two gradients, the frequency-encoding or readout, and the phase-encoding gradient. These gradients are actually just different implementations of the same process, chiefly differing only in the duration and timing of their application. To see this, it will be useful to examine the relationship between frequency and phase.

When we think of frequency and phase, we are apt to think of the familiar basic sine and cosine functions. Translated to the unit circle, these functions can be seen as a vector rotating around in a counterclockwise manner. The radian frequency is then defined as $\omega=2\pi f$. For most applications, it is convenient to define this circle in the complex plane and utilize the Euler relation. $e^{i\theta}=\cos\theta + i \sin\theta$. Now if a signal completes 2π radians in one cycle then we say that it has a frequency of one Hertz. Alternatively we can say that the vector undergoes a phase shift of 2π radians over the one second. Therefore, we can define frequency as the derivative of phase with respect to time, $d\theta/dt$. In this manner, phase and frequency are intimately related and ultimately encode space in the same manner (Wehrli, 1988).

Let us first consider the frequency-encoding gradient. After the application of the slice-selection gradient the spin system can be likened to a harp, where all the strings are vibrating at the same frequency. That is, there is no way of determining when a note is played which, if not all of the strings, contributed. The application of the frequency-

encoding gradient is then analogous to the tuning of the harp. When the harp is tuned, each string has a unique frequency, and the trained observer can then easily determine which string was played by an analysis of the emitted sound. In MRI, this tuning is accomplished by the application of the frequency-encoding gradient and the trained observer is the Fourier transform.

In contrast to slice selection, frequency encoding is performed during the signal's detection, hence the alternative name for frequency encoding, the readout gradient. The frequency -encoding gradient is applied in such a manner that it alters the resonant frequency of the spins depending on their position in the slice. This gradient is applied while the signal is being sampled by the analog-to-digital (A/D) converter and left on during the entire sampling period. The effect is that the frequency of the spins evolve in time while the points (typically some power of 2 such as 128, 256, or 512) are sampled. Upon FT a map of spatial frequencies in one dimension is obtained. To acquire spatial frequencies in the other dimension, we could simply rotate the frequency gradient, much as is done in CT or PET. However, virtually all MRI systems in use today utilize another orthogonal gradient, called the phase-encoding gradient to perform this task.

1.8.3 Phase-Encoding Gradients

The frequency-encoding gradient simply causes incremental phase shifts in spins located at different positions in our sample during the entire readout period. The phase-encoding gradient also causes a phase shift in the spins. In fact, the only difference between the two is a matter of timing, i.e. when the gradient is applied and the duration of its application. Basically, in phase-encoding a gradient is applied before the readout period which shifts the phase of the spins by some amount α, this amount is dependent upon the

amplitude of the gradient. After the phase-encoding gradient has been applied, the signal is sampled in the presence of the frequency-encoding gradient. By varying the amplitude of the phase-encoding gradient initial phase shifts of various degrees can be created. Accordingly, the main difference then between phase and frequency-encoding is that during frequency-encoding all the shifts in phase are collected at one time, while with phase-encoding a separate signal is acquired for each increase or decrease in the gradient amplitude. Accordingly, N repetitions are required to acquire N signals. In this way, it is clear than in a conventional imaging sequence the time required to collect an image is dependent upon the number of different amplitude increments used in the phase-encode process.

Once the phase-encoding gradient has been applied a suitable number of times, say 128 or 256, we obtain a rectangular grid of complex-valued data points, where by convention the x-axis represents frequency and the y-axis phase. This complex-valued dataset is referred to as *k-space*. From k-space data, we can perform a two-dimensional Fourier transform (2DFT) to obtain an image. It should be noted that the Fourier transform also results, in general, in complex-valued data. Recall that any complex number z can be written as $z=|z|e^{i\arg(z)}$, where $|z|$ is the modulus, or magnitude of the complex number and $arg(z)$ is the principle value of the angle of the vector. Typically only the moduli are used in the formation of a MRI image. This is referred to as a magnitude image. A phase image can also be formed, but this is less commonly encountered. In the upcoming discussion of fMRI time series it will be useful to keep in mind that the time series will consist of a series of moduli resulting from the 2-DFT of k-space data. Finally, it should be noted that the orthogonality of the gradients is the most efficient design. Why this is so should be

clear from a consideration of vector spaces. A *basis* in an N-dimensional (finite) vector space is a linearly-independent minimal set of vectors such that every other vector (point) in the space can be written as a unique linear combination of the basis vectors. Orthogonal bases are particulary amenable to work with, since any set of N vectors which are mutually orthogonal will constitute a basis for an N-dimensional space. Accordingly, in a 2D image, we need two basis vectors (the phase and frequency-encoding gradients). In a 3D image we would require three, but as we add more vectors we must take care that they are linearly independent to ensure that our set still constitutes a basis for the higher dimensional space. If the mutual orthogonality of the vectors is maintained, this result is guaranteed.

1.9 Pulse Sequence Design (PSD)

The construction of an image using NMR, requires a complicated interplay between the slice-selection, frequency-encoding, phase-encoding gradients, and RF pulses called a *pulse sequence*. Pulse sequences typically fall into to one of two general categories, or types: *spin echo* (SE) or *gradient echo* (GRE). Both sequences involve a "refocusing" of the spins before collecting the signal, but they differ in the way they achieve this refocusing and consequently there are differences in the sensitivity of these sequences to T2 vs T2* effects. A spin echo sequence uses a 180-degree RF pulse in addition to a gradient, while a gradient echo, or gradient-recalled echo lacks the 180-degree pulse. Refocusing here entails dephasing the spin system deliberately and recreating the conditions for a FID at a later time. To understand what is meant by *echo*, recall that the signal produced by a single RF pulse is an exponentially-decaying complex of sinusoids.

Also recall that RF pulses are characterized by how much they "flip" the net magnetization of the spin system with respect to the external field. Various *flip angles* are illustrated below.

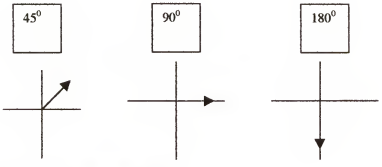

Figure 1-Flip angle illustration

Finally, remember that the oscillating magnetic field of the spin system will be most easily detected (maximum signal) when it is furthest away from the external field. In vector terms this will occur when the two vectors are orthogonal. Hence, the 90^0 flip angle will yield the largest signal. Conversely, a flip angle of 180^0 should make the spin system undetectable since it is again lined up with the external field, but in the opposite direction. It should be kept in mind that even though the 90^0 RF pulse provides maximal signal, RF pulses with flip angles less than 90 degrees are extremely useful in gradient echo pulse sequences.

Now, the FID which occurs after an RF pulse occurs too quickly for application of the gradients needed to spatially-encode the signal. However, if the initial 90^0 pulse if followed by a 180-degree pulse, the spin system can be inverted and then allowed to relax, effectively postponing the FID until such time when it can be collected after the application of the phase-encoding and frequency-encoding gradients. To understand this process in more detail, one can think of the relaxation processes as restoring the

magnetization vector from the 180-degree orientation back to equilibrium (0 degrees) in a counterclockwise fashion. As the vector approaches 90^0 a FID will be generated. In an analogy with an audio tape recording, the FID generated by a 90-degree pulse is simply playing the tape forward. The 180^0 pulse is analogous to first rewinding the tape, or more exactly playing it in reverse at the normal speed, and then allowing it to resume playing. An echo is created by reflecting the FID about the y-(amplitude) axis. The term spin echo has been used here, since a 180-degree pulse was used in the example, but as stated previously, the same process can be achieved by using gradients alone. There are, however, important differences between the refocusing achieved with a 180-degree pulse and that obtained through the use of gradients alone. To understand this effect, consider a simple analogy. Assume there is a track race, with all the runners lined up at the starting line. When the gun is fired (the 90^0 pulse) all the runners are even (in phase). However, the runners begin to disperse (dephase) because they are traveling at different speeds. In NMR terms the protons are the runners and they begin to dephase because of the processes we have described earlier. If at some time T we fire the starter's gun again (180^0 pulse) we can force the runners the runners to return to the starting line. If the runners return at the same pace they will all end up back at the starting line at exactly the same time. They will be back in phase and the signal will once again be at its maximum. In this way the signal has been recreated at a point in time where we are in a position to collect it. Additionally, something remarkable has been accomplished. Recall that T2* refers to spin-spin interaction and the additional dephasing effects of magnetic field inhomogeneity. Also, we have seen that T2 effects are random, owing to the random nature of molecular interactions. However, if we think about magnetic field inhomogneity effects, we can

assume that these are static. Therefore, we can assume that any component of dephasing caused by magnetic field inhomogeneities can be reversed by the use of an 180^0 pulse. This is, in terms of our analogy, requiring that the runners return to the starting line at the same pace with which they departed. In this way we can cancel out any static dephasing effects. This works because any static effects that occur up to a time T will be the same that occur in any time interval of length T. However, the random effects of spin-spin interaction cannot be reversed since they will not be the same from T to 2T. Now we are in a position to examine a simple spin echo pulse sequence. Consider the following chain of events: 1.) The slice selection gradient (Gz) is applied simultaneously with the RF pulse, ensuring that only the region of interest is excited. 2.) The phase-encoding gradient is applied between the 90^0 and 180^0 pulses. Recall that the amplitude of this gradient is altered before the acquisition of each line in k-space. 3.) The frequency-encoding, or read-out gradient is applied while the echo is being collected. Now, there is an important basic fact about the application of gradients. Each time a gradient is applied, the spins are dephased. By first applying a negative gradient and then applying the positive gradient, the spins are first reversed in much the same way as a 180^0 pulse. By keeping the area of the negative gradient to be ½ of the positive, the signal will re-focus, or be at its maximum, at the center of the positive gradient. By synchronizing the spin echo produced by the 180-degree pulse and the center of the positive gradient, maximal signal is obtained and perfect refocusing has been achieved. Since the goal here is not to provide an exhaustive description of pulse sequence design, it will suffice to state that gradients always produce dephasing of the spin system. Therefore, refocusing gradients are typically employed to offset the loss of signal. In practical terms, this means that a spin echo sequence actually

contains two echoes, one produced by the use of a refocusing gradient and one produced by the 180-degree pulse

As previously stated, spins can be refocused with a gradient instead of a 180^0 pulse. This is done by using a gradient with reversed polarity to cause spins precessing at a certain frequency to begin precessing at an equivalent negative frequency. In vector terms, we cause the vector to rotate at the same radian frequency but in the clockwise direction. In this way, the sample can be refocused in a way similar (but not identical) to a spin echo. The type of sequence that utilizes a gradient to perform this refocusing is called *gradient* or *gradient-recalled echo*. Although they both refocus the sample, the spin echo and gradient echo sequences are quite different in a number or respects. Firstly, gradient echo does not eliminate the effects of T2* decay. To see this, consider a spin, which inhabits a point in the gradient field that is offset from the Larmor frequency by 200 Hz. This spin may actually be precessing slightly above or below that frequency because of static magnetic field inhomogeneities, say +/- 10 Hz. If we apply a gradient reversal we have caused that spin to precess at -200 Hz for a predetermined time. However, we have not been able to compensate for the +/- 10 Hz offset because we have not completely flipped the spins as is done in a spin echo sequence. For this reason, gradient echo sequences are sensitive to the full T2* effect. Also, in contrast to spin echo sequences, gradient echo sequences are faster, because gradients can be reversed much faster than additional RF pulses can be introduced. Gradient echo sequences also often use flip angles, which are less than 90^0. Intuitively this will allow for faster imaging since the longitudinal magnetization recovers faster if it is been flipped less than 90 degrees. Faster data acquisition is the chief advantage of GRE sequences over SE Furthermore, because

gradient echo does not use additional RF pulses, the RF power deposition is much less in gradient echo sequences.

Finally, while it may appear initially that the ability of SE sequences to eliminate T2* effects is a distinct advantage over GRE sequences, this is application specific. For example, the BOLD contrast is primarily a T2* phenomenon, and while it can be observed using SE sequences, it is much easier to detect with gradient echo.

1.10 Image Weighting

MRI as an imaging modality would not be nearly as useful if different tissue types could not be distinguished based on their chemical composition. Let us consider how we can manipulate the relaxation properties of a tissue to construct an image with various degrees of contrast by varying some parameters of data acquisition.

Imagine we have two types of tissue, one with a long T2 and the other with a short T2. If we look at the course of transverse relaxation after the application of an RF pulse, we will see that the at some point the tissue with the short T2 will have already lost nearly all of its coherent precession. On the other hand, the material with the long T2 will still have a substantial transverse component. Since we are free to collect (or not to collect at all!) the signal at any point after the application of the RF pulse, we are able to use this to differentiate tissue based on their T2-T2* relaxation properties. The time interval between the excitation of the spin system and signal collection is a user-determined quantity called the *TE (time to echo)*.

Yet another parameter we can vary is the time between RF pulses. Manipulating this quantity primarily causes contrast difference in materials with different rates of T1 relaxation. Recall that T1 relaxation refers to the process by which the longitudinal

magnetization returns to equilibrium. If a spin system is allowed to completely recover before the next excitation, then number of protons eligible for resonance will be at its maximum and so will signal intensity. Conversely, if the spin system has not fully recovered, then fewer of the spins will be eligible for resonance and the resulting signal will be less intense. The time between applications of RF pulses is called the *TR (time to repetition)*.

The signal can be weighted in favor of T1 or T2 differences in the material by manipulating both the TR and TE. For example, if a long TR is used, while varying the TE, the image will be T2- or T2*-weighted. Differences in image contrast will chiefly be a result of differences in the T2-T2* properties of the tissues. Alternatively, if a short TE is used and the TR is varied, image contrast will be based on T1 differences. With a short TE, there will be no time for any substantial loss of coherence even in tissues with short T2s. However, materials with long T1s will not have recovered fully if a sufficiently long TR is not used. The resulting image is called *T1-weighted.* Finally, we can collect an image based on the number of protons of interest in the sample, i.e. the more present the stronger the signal, by using *proton-density weighting*. A proton-density weighted image is collected by using a short TE and a long TR, thereby minimizing the effects of both T2 and T1 relaxation and allowing signal intensity to be based on the number of visible spins in the different regions of the image/sample. It should be noted that it is very difficult to obtain an image which is pure T1- or T2-weighted, in fact, there will always be some T2-weighting since T2 times are typically much shorter than T1. Nonetheless, these principles have provided powerful diagnostic tools for the evaluation of normal and pathologic tissue states in living organisms.

1.11 Fast MRI

In conventional MRI, the time required to collect an image is largely determined by the TR and the number of phase-encoding steps. As a result, the time required to acquire an MRI image is considerably longer than other imaging modalities, such as ultra-sound and CT. This turns out to be one of the most serious drawbacks for MRI since long acquisition times increase the possibility that patient motion or flow-related artifacts will distort the image. In functional imaging, changes in cortical activation and the accompanying increase in cerebral blood flow occur on a time scale much shorter than a typical imaging sequence's acquisition period. Accordingly, useful information will be lost if the temporal resolution is not considerably greater than in a typical clinical acquisition. Fortunately, several methods have been developed which allow an image to be formed in the second or sub-second range. All fast imaging sequences share the fact that they acquire multiple phase-encoded lines in each excitation.

The overwhelming majority of fMRI studies to date have used one of two fast imaging sequences: *Spiral scanning* or *Echo-planar imaging*. Both are capable of acquiring all the data needed for a 2-D image in one excitation, or a single *shot*. However, in practice it is more common to encounter the spiral sequence in a 2-shot or 4-shot implementation. The actual fMRI data used in the current study were acquired using single-shot EPI.

1.11.1 Echo-Planar Imaging

Currently, there are several different varieties of EPI used in a number of different applications. EPI can be implemented in a single-shot or mutli-shot PSD. Furthermore, EPI can be utilized as a gradient echo, spin echo, or hybrid (asymmetric spin echo)

sequence. Images in the current study were collected with a single-shot gradient echo technique and therefore this discussion will be limited to that particular PSD.

Single-shot gradient-echo EPI refers to sequence in which all the lines in K-space are filled by multiple gradient reversals producing multiple gradient echoes after a single RF pulse. As in the standard gradient echo sequence, each echo requires that the frequency-encoding gradient be applied with a refocusing side lobe of opposite polarity. In single-shot EPI, the frequency-encoding gradient must be rapidly reversed between maximum positive to maximum negative polarity once for each echo. For 128 points this requires 256 lobes during a single T2* decay. As a result, EPI places tremendous demands on the gradients with respect to both strength and rise/fall time.

In the earlier EPI methods, the phase-encoding gradient was kept on continuously during the acquisition. This method led to k-space being covered in a zig-zag method. Unfortunately, this method resulted in some artifacts during the application of the 2-D Fourier transform. A solution to this problem was offered by the introduction of *blipped EPI*. Blipped EPI refers to the application of the phase-encoding gradient for a brief period when then frequency-encoding gradient is zero, i.e. when the trajectory in K-space is at one end of each frequency-encoded line. This type of trajectory has proven much more easily handled by the Fourier transform and as a result, it is the most common phase-encoding technique used today. The following diagram depicts a simplified PSD for a single-shot gradient-echo EPI.

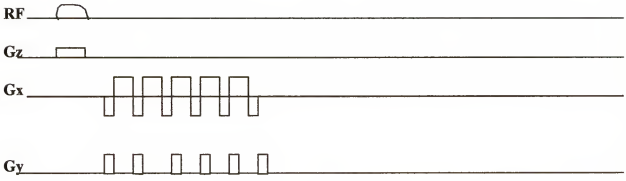

Figure 2-Pulse Sequence Diagram

Single-shot EPI has several advantages and disadvantages with respect to multi-shot sequences. A distinct advantage of single-shot EPI is its ability to produce a 2-D image in the sub-second range. This makes it ideally suited to cardiac and respiratory imaging where motion artifacts pose serious problems for more conventional sequences. Additionally, in fMRI, where speed is critical, single-shot EPI has proven quite effective. Some of the disadvantages of single-shot EPI include: 1.) The possible induction of currents in the subject due to rapid switching of powerful gradients, and 2.) The potential for phase-error propogation, caused by slight variations in resonant frequencies. This latter problem is exacerbated by intrinsic non-uniformities in B_0.

1.12 Diffusion and Perfusion in MRI

Random molecular motion was encountered in the discussion about T1 and T2 relaxation. The fact is that random molecular, or Brownian motion is an inherent part of the physical universe. Indeed, only at temperatures near absolute zero do we begin to see what the universe would be like if there were no random molecular motion. When we think of diffusion, we immediately think of the motion of molecules down a concentration gradient, e.g. the way a drop of ink will spread out in a glass of water. However, diffusion

is governed by random processes, which take place in the absence of any concentration gradient. In fact, if the course of an individual molecule in a free medium could be followed, it could be shown that in most cases its course is only predictable in a probabilistic sense. In probability theory, these processes are called random walks. Diffusive processes are ubiquitous in living organisms where they serve many important functions, including the supply of nutrient molecules from the capillary beds to cells that depend on those nutrients. In the brain, these diffusable substrates are primarily glucose and oxygen. NMR has proven quite useful for the study of these diffusive processes because it is possible to label molecules by manipulating the magnetization of the hydrogen proton without interfering with the proton's random walk . Presently, diffusion imaging has become an extremely important application of NMR in many areas including the evaluation of edematous tumors, as an early indicator of ischemia, and in the investigation of demyelinating diseases (Demaerel et al., 1999; Miller et al., 1998). As in fMRI and conventional anatomical imaging, diffusion techniques most often target hydrogen protons for the same reasons discussed previously.

As previously mentioned, diffusion is typically described as the movement of molecules (the solute) in some medium (the solvent) from an area where the solute is highly concentrated to an area of lower concentration. This is illustrated in classical physics by the development of Fick's law, which parallels the equation for heat conduction. Fick's law can be stated in a number of different but equivalent ways. For example, we can say that the mass m of solute that diffuses in a time t through a solvent contained in a channel of length L and cross-sectional area A is given by

$$m = \mathrm{D}A\,\Delta \mathrm{C}t/L$$

where ΔC is the concentration difference between ends of the channel and D is the diffusion constant analogous to the term for thermal conductivity in the heat conduction equation. Alternatively, we could use vector gradient notation, and ignoring the contribution of the physical medium in which the concentration is located, and describe it as

$$J = -D\nabla C$$

where ∇C refers to the concentration difference (or gradient) and J is mass per unit time, or particle flux.. The development of Fick's law, as with most physical models for natural phenomena, relied on substantial simplifications and idealizations. Still, it can be shown that even in a complex medium such as the human brain, diffusive processes obey Fick's law subject to two important modifications. The most important modification results from the number of obstacles to free movement of water in the human brain. This fact requires that the diffusion coefficient D in Fick's law be allowed to assume different values subject to environmental constraints. In many media, hydrogen protons are not free to wander around in any direction with equal probability as would be expected in a strict random walk model. Instead, their movement must be modified to accommodate the presence of fibers, intracellular organelles, and macromolecules. As might be expected, studies of diffusion in the human brain have found marked differences between the diffusion coefficient of CSF, white matter, and gray matter. For example, *in vivo* studies have demonstrated that the diffusion of water in CSF is similar to that in free water with $D \approx 2.95 \times 10^{-3}$. In gray matter diffusion is approximately 2.5 times lower than in CSF. The picture of diffusion in white matter appears to be much more complicated, with the diffusion coefficient dependent upon the orientation of the myelin tracts. In fact, while

diffusion in gray matter and CSF basically occurs with equal probability in all directions, *isotropy*, diffusion in white matter is restricted to particular directions. This is known as *anisotropy*.

The utility of diffusion imaging as a clinical and diagnostic tool relies on the fact that many pathologic changes in the brain affect the diffusion coefficient. For example, when neurons are deprived of oxygen and glucose, as in ischemia, cells begin to absorb NaCl and consequently swell. This swelling serves to further inhibit the free movement of water molecules and these changes are detectable on diffusion-weighted images before they are evident on more conventional structural images.

Perfusion in living organisms generally refers to the delivery of nutrients to the cell. In the human brain, these nutrients are glucose and oxygen. Glucose is the raw material for the manufacture of energy-providing ATP and the catabolism of oxygen is the most efficient means of production. Oxygen and glucose in the human brain are diffusable substrates and as such are delivered to cells from the cerebral microvasculature via diffusive processes. Accordingly, at a very early stage it was recognized that the concepts or diffusion NMR could be extended to evaluate blood microcirculation and perfusion. True to this historical connection, *arterial water labelling* is a endogenous technique which measures perfusion by magnetically labelling water protons in arterial blood. Other techniques utilize exogenous tracers in ways analogous to PET and SPECT. Finally, the blood oxygenation level dependent contrast (BOLD) is a way of indirectly measuring perfusion that relies on rapidly changing magnetic susceptibility differences. It is this last technique which has prompted the most recent and most profound incursion of MR into the field of functional neuroimaging.

1.12.1 The BOLD Contrast

Since the last decade of the 19th century, it has been known that local variations in neuronal activity lead to local increases in the vascular supply. Throughout this century, a number of techniques have been developed to measure these changes in the vascular supply in an attempt to characterize the functional activity of the brain. The majority of the early attempts were based upon the perfusion-limited distribution of non-metabolized and freely diffusable tracers. These methods are still in use today, forming the basis of PET and SPECT. However, developments in the late 1980's and early 90's indicated that MR methods relying on an intrinsic contrast mechanism could also be used to follow rapid changes in the vascular supply. Whereas conventional proton density, T1 and T2 imaging had not proven sensitive to rapid changes in the physiologic state of tissue, Moseley et al. 1989 demonstrated that within 30 minutes of the onset of ischemia, the diffusion coefficient of the brain started to decrease. Work on the development of MR as a functional technique began to attract serious attention.

In the early 1990's it was demonstrated that MR images could be made sensitive to the rapidly-changing magnetic susceptibilities of voxels brought about by alterations in local cerebral blood flow (Ogawa et al., 1990; Kwong et al., 1992). This became the foundation of the currently most widely-used method of fMRI, the BOLD contrast. The basis of the BOLD contrast is the different magnetic susceptibility of hemglobin depending on its oxygenation state. Hemoglobin is the protein which carries oxygen to the erythrocytes. Hemoglobin is comprised of a polypeptide chain, globin, and the heme group, which contains iron. In its deoxygenated state, three of the four outer electrons on the iron atom are unpaired. In oxyhemoglobin, an oxygen atom binds with the iron in the heme group

and reduces the magnetic susceptibility effects of the unpaired electrons. A paramagnetic substance, as we have stated, is one in which some of its electrons line up with an applied magnetic field, producing their own additive fields in the process. Diamagnetic materials counteract the effect of the magnetic field, decreasing its effect inside of them (Wehrli, 1988). The additive fields produced in a paramagnetic material cause a reduction in the NMR signal because they contribute to the dephasing of the spin system. In more technical terms, the presence of a paramagnetic agent creates a susceptibility difference between the blood and the surrounding tissue. In T2 or T2* -weighted images this susceptibility difference shows up as a decrease in signal intensity, since the spread in phase angle within a voxel results in some degree of phase cancellation. This fact was demonstrated convincingly by Ogawa and his colleagues when they imaged the brain of a rat at 7T while the rat's O_2 intake was regulated. They found the signal intensity was much larger on T2- and T2*-weighted images when the rat breathed 100% O_2 as opposed to 20%. They also found that the effect, while present, was reduced on spin echo as opposed to gradient echo images. This fact provided support for the conclusion that the phenomenon was the result of magnetic susceptibility differences. Kwong et al., 1992 demonstrated the BOLD contrast on GE images in humans using visual stimulation and subsequently these results have been replicated by various groups with a number of sensory and cognitive tasks and at a number of different field strengths.

The most widely used model for the BOLD contrast is a cylinder, which has a magnetic susceptibility different from the surrounding fluid. This is similar to the models used with paramagnetic contrast agents such as gadolinium, but the susceptibility difference induced by deoxyhemoglobin is an order of magnitude less. For example, for fully deoxygenated

blood with a hematocrit of 40%, the susceptiblity difference is around 0.07 ppm, which corresponds roughly to a gadolinium concentration of 3 mmol/L or about 0.5% of commercially available solution. Accordingly, the signal changes seen in BOLD studies are much lower than seen with gadolinium, typically on the order of a few percent, although areas surrounding major vessels can exhibit percent changes on the order of 15-20%. These small signal changes have prompted many to push for higher field strengths. Functional imaging studies conducted at 4T have reported signal changes much larger than those seen at the standard clinical imaging field, 1.5T, but it is not clear whether a higher field strength gives a clear advantage, or at least one substantial enough to warrant the cost.

1.12.2 Arterial Tagging Methods

Many studies have demonstrated a linear relationship between increases in regional cerebral blood flow (rCBF) and regional consumption of glucose (rCMRgl) (Baron et al., 1984). However, the relationship between rCBF and regional oxygen consumption ($rCMRO_2$) is less clear. The fact that the BOLD contrast is actually detecting changes at least two steps removed from the actual neuronal activity has led some in the fMRI community to advocate methods which measure perfusion directly. Currently, these methods are all based on the concept of *arterial tagging*. Arterial tagging is a subtraction technique similar to the O^{15} method for measuring perfusion in PET. In NMR, relaxation plays the role of radioactivity in PET. In arterial tagging methods a the spins in a slice adjacent to the slice of interest are inverted, then as those spins flow or diffuse into the slice of interest they will still be inverted provided the transit time is sufficiently short. Images are collected both with and without the inversion pulses and signal intensity

differences are held to be proportional to the amount of blood delivered. Because quanitative methods in fMRI are still in their infancy, most comparative studies between BOLD and arterial-tagging techniques have been qualitative. Thus far, the two techniques have demonstrated similar results. Furthermore, even though arterial-tagging may indeed represent one step closer to the actual neuronal activity, there are potential problems with its application. These include incomplete inversion of the arterial blood and relaxation during transit. Also, signal changes obtained with arterial-tagging techniques tend to be smaller than those seen in BOLD studies. This coupled with the fact that arterial-tagging methods are harder to implement continue to make their implementation much less frequent than the BOLD technique.

1.13 Brain vs. Vein Debate

A major concern when using BOLD fMRI to map human cognition is how focal are the signal changes to the actual underlying electrophysiology. Because the BOLD contrast has only recently been demonstrated in an anesthetized animal using sensory stimulation, adequate empirical evidence from single-unit studies in animals compared with BOLD maps is still lacking. Recent work has demonstrated excellent agreement between local increases in cerebral blood flow and the electrophysiologic response in the whisker barrel of the rat. However, this relationship is certainly more complex than this one study suggests. Further, even if the coupling between the underlying neurophysiology and local increases in cerebral blood flow is quite tight, there is still the matter of how accurate we are in detecting only those capillary level changes as opposed to more global blood flow and blood volume changes in the brain. This is of particular concern when using T2* methods in detecting the BOLD contrast. Several studies have shown that GE images are

very sensitive to changes in the concentration of venous deoxyhemoglobin in large vessels. Ideally we would want to suppress these large vessels signal increases because they are likely to be more diffuse than changes at the capillary level. Spin echo, or hybrid sequences such as ASE have been shown to be more sensitive to small vessel signal changes. However, using spin echo sequences has been eschewed by many groups because of signal attenuation. Also, as the push continues for higher and higher fields, there is also concern over the RF power deposition found in SE sequences. This increased power deposition is the result of the additional 180° pulses.

This chapter has served as a general introduction to magnetic resonance with an emphasis on the creation of a 2D-image from resonance effects induced in a spin system. For our purposes, the most important point is that magnetic resonance imaging is a highly flexible tool, which allows one to uncover a wealth of information about a spin system by differentially weighting the image acquisition. The most salient application here is the ability of MRI to differentiate between voxels based on the concentration of deoxygenated vs. oxygenated hemoglobin. The next chapter will address the neurophysiological basis of this contrast and examine what is known about its coupling to neuronal activation.

CHAPTER 2
FUNCTIONAL IMAGING-PHYSIOLOGICAL BASIS

The principle application of functional magnetic resonance imaging (fMRI) is to aid in understanding how the human brain processes information. Information processing must be understood here in the broadest sense, since fMRI is currently being used to investigate a wide variety of sensory, motor and cognitive processes. However, despite the vast differences in experimental designs, all fMRI studies at least tacitly assume that their data represent *brain activation*. This implies that we understand what is meant by brain activation. It further implies we can claim that signal changes based on differences in regional cerebral blood flow (rCBF), cerebral blood volume (CBV), and blood oxygenation represent neuronal activation. As it turns out, these are not simple claims to support. To answer them requires we understand the relevant neurophysiology, statistical theory, and NMR methods. In previous chapters we have devoted considerable attention to NMR methods and the basis of the BOLD contrast. Here we will undertake to examine in some detail the neurophysiological basis of rCBF changes and how those changes are coupled to electro-chemical changes in the brain. Furthermore, we will examine how signals based on the rCBF changes are analyzed and corresponding inferences drawn.

2.1 Brain Activation

Imagine that you are sitting in a restaurant when an unfortunate waiter suddenly drops an entire plate of dishes to your right. Naturally, you immediately turn your head to that

side to see what happened. Your reaction seems trivial, but in fact if we examine it from a neural information-processing perspective, it is extremely complex. First, your brain had to undertake a complex analysis of the acoustic properties of the sound to determine its probable source. Secondly, a complicated cascade of neuronal firings had to occur before you turned your head in the direction of the sound source. If we could examine these processes with infinite temporal and spatial precision we would see that there were a number of changes in brain activity evoked by the sound. These fast, transient changes are a prime example of what is understood to be *brain activation* in fMRI. In other words, we are not concerned with activity associated with brain *states* such as sleep, being awake, or being unconscious, or transitions between such states. Rather we are interested in activity evoked on a much shorter time scale, typically in response to some stimulus. The activity may or may not be evoked by primary sensory stimulation. It could for example be the result of trying to remember something, imagining some scenario, or performing a mental calculation. Furthermore, brain activation in fMRI implies that a large population of neurons is activated. This is a necessity given that the NMR signals generated by the BOLD contrast are relatively weak.

It is well established that the basis of all information processing in the brain is the result of biochemical changes within neurons. These biochemical changes are in turn communicated to neighboring neurons over synapses. As we will see, the effect on the neighboring neuron may take one of several different forms and is usually the result of a complicated integration of activity.

2.2 Cellular Components of the Mammalian Brain

The mammalian brain is comprised of two general classes of cells, neurons and neuroglia. The neuroglia, or "nerve glue" are the most numerous, constituting about 80-90% of the cells in the CNS. Neuroglia can be divided into two broad categories: macroglia and microglia. Macroglia consist of astrocytes, oligodendrocytes, and ependymal cells and are embryologically ectodermal. The microglia are thought to originate from the mesoderm. The glial cells differ from neurons in the following ways: 1.) glia cells retain the ability to divide in the mature CNS, 2.) they do not form synapses, and 3.) they do not transmit action potentials. Of the three types of macroglia, the astrocytes are the most abundant in the mammalian CNS. In addition to comprising the major structural support for the CNS, they also are implicated in the regulation of cerebral blood flow.

Neurons can only be described in a general sense since the morphology of an individual neuron is determined in large part by genetics and its role in processing information. However, we can say that all neurons consist of a cell body, or *soma*, and cytoplasmic processes extending from the cell body. The cell body contains the nucleus and cytoplasm, which contains the organelles needed for metabolic activity within the neuron. The cytoplasmic processes extending out from the cell body are called *dendrites* and the *axon*. Each neuron only has one axon, while the dendrites may branch out from the cell body in elaborate patterns called *arborizations*. The axons and dendrites form the connection between neighboring neurons which is the basis of the transmission of electrical activity. These connections are referred to as *synapses*.

2.3 The Synapse

Synapses are specialized connections, or junctions between neurons and other cells. Neurons communicate not only with other neurons, but also with effector cells such as skeletal muscle fibers. There are two types of synapses in the mammalian brain: electrical and chemical synapses (Roland, 1993). The most abundant of these is the chemical synapse. A chemical synapse consists of a presynaptic terminal, a small extracellular space called the synaptic cleft, and a postsynaptic thickening. The synaptic cleft is typically about 20 to 40 nm wide. Initially thought to be an empty space, it is now believed that the synaptic cleft may in fact contain some type of carbohydrate material. To see how a chemical synapse functions, let us review some of the basic mechanisms of neuronal activation.

The typical mammalian CNS neuron has a resting electrical potential of about -70 mV. This is chiefly a result of negatively-charged groups on protein molecules in the cytoplasm. Furthermore, the neuron membrane is typically impermeable, isolating ions in the extracellular fluid from those inside the cells. Mammalian extracellular fluid has a resting potential around 0 mV. Therefore, the cell is resting about 70 mV more negative than the extracellular fluid and an electrical gradient has been created between the two. Events we will discuss in some detail can cause ion channels in the cell membrane to open allowing an influx of positive Na^+ ions to enter the cell. This causes the cell to depolarize which in turn causes the neuron to generate an *action potential* which propagates along the axon at a constant speed. In a chemical synapse, the action potential causes a release of neurotransmitter at the presynaptic terminal. This neurotransmitter then diffuses across the synaptic cleft and binds to receptors (glycoproteins) on the postsynaptic terminal. This

binding then can cause an *excitatory postsynaptic potential* (EPSP), an *inhibitory postsynaptic potential* (IPSP), or neither. The probability that the postsynaptic neuron will undergo a similar process and produce its own action potential is dependent on the complex spatio-temporal summation of EPSPs and IPSPs it receives from other neurons.

2.4 Neurotransmitters

Neurotransmitters may be divided into three categories depending on their chemical composition and mode of action: amino acids, amines, and peptides. The receptor molecules for these transmitters are glycoproteins and can be located on any number of different sites along the neuron including the presynaptic terminal, the soma, or the axon (Roland, 1993).

The amino acid neurotransmitters typically cause fast depolarization or hyperpolarization. Acetylcholine, which is actually an amine, sometimes acts like an amino acid transmitter in this regard. The two major excitatory amino acid transmitters in the mammalian CNS are glutamate and aspartate. They are termed excitatory because they cause ion channels on the postsynaptic neuron to open which selectively allow Na^+ and less frequently Ca^{2+} to enter causing depolarization, pushing the cell toward the generation of an action potential. There are actually four types of excitatory amino acid receptors, the most well-known of these being the NMDA receptor. Research has shown that glutamate and aspartate bind to all four types of excitatory receptors, although it is not clear what if any differences there are in their effects. It is, however, thought that glutamate is the most abundant of the two in the CNS (Roland, 1993).

The major inhibitory neurotransmitters are gamma-aminobutyric acid (GABA) and glycine. These neurotransmitters cause the neuron to open ion channels selective for Cl^-.

Of the two, most studies indicate that glycine plays the most restricted role, with its action being chiefly confined to more rostral areas of the CNS and the retina.

In contrast to amino acid transmitters, amines typically act slowly over the course of seconds or even minutes. For this reason they are often referred to as neuromodulators. Acetylcholine acts on two kinds of receptors: nicotinic and muscarinic. When acetylcholine binds to a nicotinic receptor it acts like an excitatory amino acid, while its action on the muscarinic receptor is more typical of an amine transmitter. Other well-known amine transmitters are dopamine, serotonin, norepinephrine, and histamine.

2.5 Energy Metabolism in the Neuron

The processes, which drive the machinery of the chemical synapse including neurotransmitter synthesis, release and re-uptake, are energy-dependent. A prime example of these energy-dependent processes is the Na^+-K^+ pump which restores the ionic balance between extra- and intra-cellular fluids. This pump in effect "recharges" the battery which forms the electro-chemical gradient necessary for the generation of action potentials. Although the brain constitutes only 2% of body weight, it receives some 15% of cardiac output and is responsible for 20% of total oxygen and 25% of glucose consumption (Roland, 1993). Energy in the human brain is derived nearly exclusively from the oxidative catabolism of glucose for the production of adenosine triphospate (ATP). ATP may be synthesized aerobically or anaerobically, but the former method is much more efficient. In vitro and animal studies suggest that the majority of ATP consumption in the mammalian brain goes to fuel the Na^+-K^+ pump. Astrup et al. (1981) anesthetized a dog with barbituates until they were unable to record any electrical activity from the animal's brain. This reduced oxygen and glucose consumption to 70% of resting values. Then, they added

ouabain, a blocker of the Na^{+}-K^{+} pump, and found that glucose and oxygen consumption fell to 35% of resting values. Mata et al. (1980) used an in vitro preparation of the pituitary gland, which they stimulated electrically with 10 Hz pulses. This stimulation produced a 30% increase in glucose consumption. After the addition of ouabain, the stimulation was repeated and this time glucose consumption did not increase.

Since fMRI primarily measures changes in the oxygenation state of blood and the proportion of oxygenated blood in a particular region, it is critical to understand how changes in blood oxygenation and cerebral blood flow are related to changes in glucose consumption and ultimately to increases in neuronal activity. Specifically, it is important to know how tightly increases in oxygen consumption are coupled to increases in glucose consumption, and how glucose consumption is related to neuronal activity.

Concerning the coupling between neuronal activity and energy metabolism, a number of experiments have demonstrated that an increase in neuronal activity does lead to a concomitant increase in energy metabolism. Larrabee (1958) showed that sympathetic ganglia increased their glucose and oxygen consumption when stimulated electrically. Further, he demonstrated that the rates of glucose and oxygen consumption depended on the rate of electrical stimulation. McIlwain (1954) stimulated brain slices and found increases in oxygen consumption, while Hand et al. (1978) showed that stimulating one whisker in the rat led to corresponding increases in glucose consumption in the contralateral cortical whisker barrel. Studies such as these demonstrate a convincing relationship between increases in neuronal activity and the regional consumption of glucose and oxygen. These energy-providing substrates are delivered by cerebral blood flow. Before examining how neurons are able to "signal" for an increase in the flow of

blood to provide these nutrients, it will be useful to understand some of the basic mechanisms of global and regional cerebral blood flow.

2.6 Anatomy of Cerebral Circulation

The brain is supplied by four major arteries, the two carotid and the two vertebral arteries. These four vessels converge on the circle of Willis, which functions as an equalizing distributor. In terms of quantity of blood flow, the carotids are the most important. It is thought that each carotid contributes approximately 40% to the total perfusion of the brain. The carotid arteries divide into two main branches: the anterior cerebral and middle cerebral arteries. Two smaller branches form the anterior choroidal, and posterior communicating arteries. Of these branches, the middle cerebral is the most important, supplying most of the lateral aspect of the hemisphere. Similarly, the vertebral arteries give rise to several branches including: the anterior spinal, the posterior inferior cerebellar arteries, and the multiple median perforating arteries. At the pontomedullary junction, the two vertebral arteries give rise to a number of paired vessels, such as the internal auditory arteries. These larger arteries run along the surface of the brain in the sub-arachnoid space. As they traverse the cortical surface, they give rise to smaller pial arteries which invaginate the cerebral cortex at right angles. The actual mode of this entry is a complicated matter subject to considerable debate. Blood draining away from the brain collects in large veins, or sinuses. In brain activation studies, this pooling often causes the sinuses, especially the sagittal and tranverse sinuses, to exhibit large signal changes correlated with the stimulus presentation.

2.6.1 Microvascular Organization

In brain activation studies, one is typically interested in the signal changes, which occur as a result of increased oxygenation of the cortical microvasculature. The reason for this is that while major vessels supply many different cortical areas, these smaller vessels have a much more restricted locus of activity. Accordingly, they are more likely to give a more spatially-restricted picture of cortical activation.

As previously mentioned, pial arteries penetrate the cortex at right angles. These arterioles give rise to capillaries at all cortical laminae. The distribution of capillaries in the brain is not homogenous. In fact, it was recognized several decades ago that capillary density in a region of the brain was highly correlated with the number of synapses (Dunning & Wolff, 1937). It is currently well accepted that capillary density provides a good indication of the oxidative metabolism demands of a particular cortical region. For example, studies have shown the density of capillaries is related to the density of cytochrome oxidase (Borowsky & Collins, 1989; Zheng et al., 1991). Furthermore, the relationship between capillary density and cortical activation has also been demonstrated in a number of animal-lesion studies. For example, orbital enucleation in rats at birth has been shown to dramatically reduce the capillary density in brain regions associated with the visual pathway, such as the superior colliculus and lateral geniculate

2.6.2 Regional Cerebral Blood Flow (rCBF)

As we have mentioned, the substrates required for energy metabolism in the brain are delivered via the blood. Oxygen diffuses across the blood-brain barrier, while glucose molecules are delivered via a carrier mechanism (Roland, 1993). Blood flow is defined as the amount of blood in milliliters transported per unit time through a unit mass or volume,

usually one gram. Although rates of CBF are not uniform across the brain, some generalizations can be made. First, a normal level of global blood flow in the brain is roughly 55 ml/100g/min. If we assume that CBF can increase locally depending on differences in regional metabolic demands, then this increase in regional cerebral blood flow (rCBF) can be accomplished in one of two ways: 1.) by letting a larger amount of blood flow through a given volume of tissue with unaltered velocity, or 2.) by letting an unaltered amount of blood flow through with greater velocity. It appears the former is the most important mechanism. In support of this, Tyson et al. (1987) and Kushinsky et al. (1987) found that the number of open capillaries was linearly related to rCBF. These studies support the idea that increases in rCBF are actually increases in the amount of blood flowing through an area and that increase is brought about by a dilation of the capillaries.

The picture emerging thus far is that an increase in neuronal activity requires an increase in glucose consumption for ATP production, which in turn requires an increase in oxygen consumption. Recall that the fuel for these processes is supplied by the blood, therefore an increase in activity requires an increase in the local blood supply. This increase is accomplished by a dilation of capilllaries in the region of increased activity. The experimental evidence emerging over the last few decades supports the idea that the brain has some mechanism for signaling an increased need for oxygen and glucose based on local activity. What mechanisms are responsible for this signaling is a matter of intense study and debate.

2.6.2.1 Nitric oxide

Furchgott and Zawadzki (1980) observed that the stimulation of endothelial cells caused the release of a diffusable factor that relaxed blood vessels under certain conditions. Subsequently, this factor was identified as nitric oxide (NO). It has been shown that excitatory amino acids cause the production and release of NO in cerebellar slices. Furthermore, the short biological half-life of NO makes it an ideal regulator of rCBF. However, objections have been raised because of the non-uniform distribution of NO synthetase in the brain. Recently, another form of NO synthetase has been discovered, so ultimately the question of uneven distribution remains unanswered (Roland, 1993).

2.6.2.2 Potassium

The typical extracellular concentration of K^+ is about 3.3 mM, although it can increase to about 7 mM during neuronal activation. Increases in K^+ have been shown to dilate small vessels by as much as 50%. The putative link between increases in extracellular potassium as the result of neuronal firing and vasodilation are the astrocytes. Recall that astrocytes are a type of macroglia. Astrocytes may be further sub-divided into protoplasmic and fibrous astrocytes. Both types have processes call *end feet*, which contact neurons and cerebral arterioles. These processes are rich in mitochondria and have uptake mechanisms for a number of neurotransmitters including, GABA, glutamate, and aspartate. They also demonstrate an interesting property called K^+ *siphoning*, which many believe is implicated directly in the regulation of cerebral blood flow. It has been shown that astrocytic end feet processes are permeable to potassium. It has also been demonstrated that extracellular potassium is siphoned from the extracellular space by these astrocytic processes and then

diffuses through the cell where it is deposited on the contacted arteriole. The rapidity of this process makes it an ideal candidate for the regulation of rCBF.

The major criticism raised against potassium has come from studies on epileptic seizures where rCBF values often remain elevated for minutes, while extracellular potassium returns to normal values in under 10 seconds

2.6.2.3 Adenosine

Adenosine is known to be a potent vasodilator. The basic idea is as follows: adenosine is released into the extracellular space by neurons and glia when neuronal activity increases. In vitro and in vivo experiments have demonstrated that electrical stimulation of the brain increases the release of adenosine (Phillips et al., 1980). Adenosine has also been shown to dilate cortical pial vessels in the dog, cat, and rat when applied topically, or by perivascular microapplication (Berne et al., 1974; Dora, 1985).

It is most likely the case that no single agent will emerge as the sole regulator of rCBF in the brain. It does, however, appear that rCBF is closely coupled to regional cerebral glucose consumption (rCMRgl). Reivich et al. (1975) found that rCBF and rCMRgl were perfectly correlated in the rat cerebral cortex. Furthermore, these processes seem to be coupled down to the level of the cortical column (Ginsberg et al., 1987).

In summary, the processes by which the neuronal machinery of the human brain operate are energy-dependent (metabolic). The required energy is almost exclusively produced by the oxidation of glucose in a series of enzymatic reactions called glycolysis. The necessary oxygen for this process is supplied by the vascular system and in particular by erythrocytes (red blood cells). The erythrocytes accomplish this task with the aid of a large protein molecule called hemoglobin. Hemoglobin consists of four polypeptide chains

(globin) and a heme group which contains an iron molecule. This iron molecule can either be bound (as in oxyhemoglobin) or unbound as in deoxyhemoglobin. Recall that the binding of the iron molecule leads to a reduction of the magnetic effect from the unpaired iron electrons and accordingly we are led to the following basic model of the neuronal basis of the BOLD contrast:

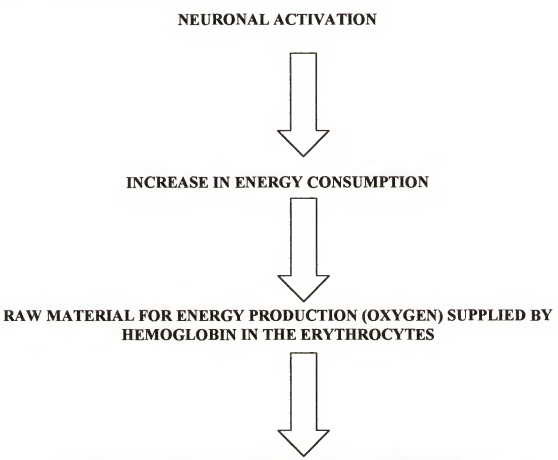
NEURONAL ACTIVATION
INCREASE IN ENERGY CONSUMPTION
RAW MATERIAL FOR ENERGY PRODUCTION (OXYGEN) SUPPLIED BY HEMOGLOBIN IN THE ERYTHROCYTES
OXYGEN BOUND TO IRON MOLECULE IN THE HEME GROUP OF HEMOGLOBIN CAUSES A DETECTABLE INCREASE IN THE LOCAL MR SIGNAL ON T2*- AND T2-WEIGHTED IMAGES

Figure 3-Neuronal basis of the BOLD contrast

In chapter one, we established the ability of magnetic resonance to detect local increases in oxygenated hemoglobin. In this chapter, the link between those increases and

neuronal activation was discussed. Armed with that knowledge, the discussion now turns to the theoretical foundations for the analysis of the BOLD signal changes in MRI.

CHAPTER 3
TIME SERIES IN FMRI

We have already alluded to the fact that fMRI is an important emerging tool in the investigation of information processing in the human brain. Since its introduction in the early 1990's, fMRI has been employed in the study of a wide-range of processes from primary sensory stimulation to working memory and emotion in both normal and clinical populations (Price et al., 1992; Binder et al., 1994; DeYoe et al., 1992; Demb et al., 1998; Perlstein et al., 1998; Lang et al., 1998). At least implicit in the conclusions of these studies is that a spatially-distributed map of thresholded significance values represents an accurate picture of the underlying neuronal activation. It is further often assumed that the activation state of a given pixel is a binary decision, i.e. there are no gradations of activated or unactivated. If a given pixel achieves the thresholded value for the chosen test statistic, it is considered activated. Otherwise it is designated unactivated and the underlying neuroanatomy is not considered involved in the task. Typically, the methods used to obtain these threshold values and the values themselves vary from study to study. Some of the most common methods currently used in fMRI are the student's T-test, the non-parametric version of the T-test, the Kolmogorov-Smirnov (K-S) test, and cross-correlation in the time domain with a reference vector (or family of reference vectors). To date, little work has been done comparing these data analysis methods on real and or

simulated data. Yet, the development of fMRI as a reliable experimental and clinical tool requires the design and implementation of appropriate and theoretically-sound approaches to data analysis. Methodological studies on data analysis must be conducted in order to determine which method(s) yield the most satisfactory results in terms of sensitivity and specificity in signal detection. This study proposes to take an important first step in that process by comparing two data analysis methods in simulated and actual fMRI data. The experimental data was obtained in a visual-stimulation paradigm at 3-Tesla. The simulated data was designed to closely approximate conditions present in actual fMRI data. The methods employed in the analysis of the real and simulated data will be multitaper spectral estimation (Thomson, 1982; Percival & Walden, 1993; Mann & Lees, 1996) and time-domain cross-correlation (Demb et al., 1998; Cox et al., 1995; Bandettini et al., 1993; DeYoe et al., 1992). Both methods are univariate procedures and assume that the data has been obtained in an experimental paradigm with a periodic forcing stimulus.

The following sections contain:

1.) Descriptions of fMRI time series with a brief exposition of the theory of stochastic processes.

2.) Description of the time-domain cross correlation method; including the derivation of its sampling properties.

3.) Theoretical background of non-parametric spectral estimation.

4.) Problems in non-parametric spectral estimation.

5.) A description of discrete prolate spheroidal sequences-their origin and utility in spectral estimation.

6.) An introduction to non-parametric multitaper spectral estimation.

In a typical fMRI experiment the intensity of a given pixel is represented as a function of time. The intensity of each pixel is typically sampled between 60 and 100 times during one experimental run and matrices of 64 x 64 and 128 x 128 of such time series are common place. The actual values of the samples are magnitudes (moduli of complex numbers) resulting from a two-dimensional discrete Fourier transform (2-DFT) of complex-valued data. Because pixel intensities are affected by a complexity of factors, we cannot model these functions as deterministic. Rather, they are an example of what is known as a random, or stochastic process. Indeed, the data obtained in fMRI are realizations of a stochastic process with complex spatio-temporal properties. The sources of the signal fluctuations are both biological and instrumental in nature. The major sources of biological noise are heartbeat, CSF flow, respiration, and low-frequency vasomotor oscillations, as well as noise arising from random molecular motion in the sample (Mitra et al., 1997). Sources of instrumental noise include: resistance in the imaging coil and transmission line, scanner drift, and computational artifacts. The fact that we cannot determine with certainty the form of these time series requires that the theory of stochastic processes be employed in their analysis.

3.1 Stochastic Processes

A stochastic process is one in which the observed values, $x_1, x_2, \ldots, x_n$ are realizations of corresponding random variables (RV), $X_1, X_2, \ldots, X_n$. A RV is a mapping from the possible outcomes (or sample space) of an experiment to the real line, the complex plane, or m-dimensional Euclidean space depending on the nature of the RV. If the sample space is the set of observable pixel intensities, then for any fixed t we can define a RV, X_t. At that point, t, X will have a probability distribution, e.g. Gaussian. This is true for each X in

the time series. In fMRI, as in time series analysis in general, we are never interested in what occurs at a fixed t, rather we are concerned with the relationship among various RVs $\{X_t,\ t \in T\}$ where t is an element of some index set, T. In most cases, T represents time, although this is not necessarily so. In a discrete time series, t is unitless, while in other applications t may represent some spatial variable. At this point, it will be convenient to introduce some notational conventions to aid the discussion:

1.) X(t) refers to a continuous-parameter random process whose component at time t is X(t). In this case t has physically-meaningful units, such as seconds.

2.) X_t refers to a discrete-parameter random process whose component indexed by t is X_t. When the index set for a random process is not explicitly stated it is assumed that it is the set of integers for a discrete-parameter process and the real line for a continuous-parameter process.

Now consider a continuous-parameter process, X(t). Recall that for a fixed t, X(t) is a RV. This implies that X(t) has an associated cumulative probability distribution function (cpdf) given by $F(x)=P[X(t) \leq x]$, where P[x] is the probability that x occurs. If they exist, the first moment and second central moments of X(t) are given by:

$$E[X(t)]=\mu_t \qquad (1)$$

$$E[(X(t)-\mu_t)^2]=\sigma^2_t \qquad (2)$$

Note that, in general, these moments depend on the index, t. Since X(t) is a random process we are inevitably interested in the relationships between the various cpdfs that constitute the entire process. This entails specifying multi-dimensional cpdfs. In the case of a continuous-parameter stochastic process this would require that we specify an infinite-dimensional cpdf, which is a mathematically intractable problem. Even in the case of a

discrete-parameter process, it is clear that this situation requires simplification. Therefore, in practice one only considers a special sub-class of stochastic processes called *wide-sense stationary* processes. In general, stationarity means that the probabilistic structure of a stochastic process is invariant under a shift in the index set. In other words, we require that $P[X_{t1} \le x_1, X_{t2} \le x_2, \ldots X_{tn} \le x_n] = P[X_{t1+\iota} \le x_1, X_{t2+\iota} \le x_2, \ldots X_{tn+\iota} \le x_n]$ where ι represents some shift in the index set. This type of stationarity, or what is termed *strict stationarity*, does not simplify our task, since we must still specify an N-dimensional cpdf for prohibitive values of N. Wide-sense stationarity, or *second-order stationarity,* requires only that the moments of order 1 and 2 are independent of the index. So, we have:

$$E[X_t]=\mu \qquad (3)$$

$$E[X^2_t]=\mu' \qquad (4)$$

Since the first and second moments are independent of the index, the variance is also independent of the index, and the covariance can be shown to depend only on the absolute difference between any two given indices $|t_2-t_1|$. In general, wide-sense stationarity does not imply strict stationarity. An important exception to this is the Gaussian distribution. If we can assume that the process $X_1, X_2, \ldots X_n$ is multivariate Gaussian then it can be shown that its distribution is completely specified by its first and second moments. Recall in the univariate case that if one knows the mean and variance of a Gaussian distribution then that distribution may be specified exactly. This result extends naturally to the multivariate case. The assumption of Gaussianity is most often made for the noise distribution. This assumption can often be justified under a central limit theorem argument since the noise arises from a number of small independent influences. Also, the resulting magnitudes (values of the time series) are the result of a Fourier transform, which in the

discrete-time case is a sum of many terms. In cases where wide-sense stationarity is not justified, it may be possible to consider the process as locally (wide-sense) stationary to simplify the analysis. In other words, if our attention is restricted to a short segment of the process at any given time, we may consider the process as stationary. This approach is common in the study of human speech, which exhibits marked non-stationarity, but may be considered as locally stationary for purposes of analysis. The theory discussed here and ultimately applied to fMRI data rests on the assumption that the process may be considered at least locally stationary. Hence, from now on, the term stationary shall be taken to mean wide-sense stationary.

3.2 Time/Frequency Analysis

The time series data obtained in an fMRI experiment can be analyzed in either the time or frequency domains, or in some time-frequency space. To date, strictly time-domain based analysis has been the most popular, although there are several potential advantages to frequency-domain based analyses (Percival & Walden, 1993). One major advantage with frequency-based approaches lies in the fact that adjacent values in the frequency domain can often be viewed as uncorrelated, greatly facilitating statistical analysis. This is not true in the time domain. Before discussing these approaches at length, let us first establish some basic results, which will aid our discussion.

- A signal is said to have finite energy if it belongs to the set of square-integrable functions, L^2, or l^2 in the discrete case. Alternatively, this is equal to saying:

$$\int_{-\infty}^{\infty} |f(t)|^2 \, dt < \infty \quad (5)$$

for the continuous-parameter case and

$$\sum_{t=-\infty}^{\infty} |f|^2 < \infty \qquad (6)$$

in the case of a discrete-parameter signal.

It can be assumed that all real world deterministic signals belong to the class of L^2 or l^2 functions.

- If a signal has finite energy then it will possess a Fourier transform given by

$$G(f) = \int_{-\infty}^{\infty} g(t) e^{-i2\pi ft} dt \quad (7)$$

and an inverse Fourier transform given by

$$g(t) = \int_{-\infty}^{\infty} G(f) e^{i2\pi ft} df \quad (8)$$

For a discrete-time sequence g_t, the Fourier transform is a periodic function of frequency given by

$$G_T(f) = \Delta t \sum_{t=-\infty}^{\infty} g_t e^{-i2\pi ft\Delta t} \quad (9)$$

where Δt refers to the sampling period and t here is unitless. The inverse Fourier transform is given by

$$g_t = \int_{-1/2\Delta t}^{1/2\Delta t} G_T(f) e^{i2\pi ft\Delta t} df \quad (10)$$

It is simple to prove that $G_p(f)$ is a periodic function of frequency with period $1/\Delta t$. Specifically, we want to show that $G_p(f+1/\Delta t) = G_p(f)$. Rewriting $G_p(f)$ in terms of its Fourier transforms

$$G_T(f + 1/\Delta t) = \Delta t \sum_{t=-\infty}^{\infty} g_t e^{-i2\pi (f+1/\Delta t)\Delta t} \quad (11)$$

which is equal to

$$\Delta t \sum_{t=-\infty}^{\infty} g_t e^{-i2\pi f t \Delta t} e^{-i2\pi} \quad (12)$$

recognizing that $e^{-i2\pi}$ is equal to 1 completes the proof. Note that if g_t represents a sampled version of a the continuous-time signal *g(t)*, then the spectrum of g_t may be quite different from that of *g(t)*. Specifically, the spectrum of *g(t)* is not likely to be a periodic function of frequency. An important theorem known as the *Shannon-Nyquist* theorem states that if a continuous time sequence is sampled at a rate at least twice the highest frequency present, then the spectra of *g(t)* and g_t will be identical over [-1/2Δt. 1/2Δt], even though the spectrum of *g(t)* may not be periodic, if we restrict our attention to an interval of length *Δt*, then the two spectra will be the same. However, in fMRI, typical sampling intervals are on the order of 1-3 seconds. These sampling rates are not high enough to satisfy the *Shannon-Nyquist* theorem if one assumes that respiration and cardiac frequencies are present in fMRI data. In fact there is convincing data available that suggests that both cardiac and respiration frequencies as well as frequencies from vasomotor oscillations on the order of 1 Hz are present in adequately sampled fMRI data (Mitra, 1997). The import of this for any frequency domain analysis is that there may be potential confounds in inadequately sampled data. resulting from a phenomenon known as *aliasing*. To demonstrate why aliasing occurs consider that an L^2 continuous-time signal has a Fourier transform given by

$$G(f) = \int_{-\infty}^{\infty} g(t) e^{-i2\pi f t} dt$$

with the inverse transform given by

$$g(t) = \int_{-\infty}^{\infty} G(f) e^{-i2\pi ft} df$$

Now, recall that the inverse transform of the sampled sequence g_t is given by

$$g_t = \int_{-1/2\Delta t}^{1/2\Delta t} G_p(f) e^{i2\pi ft\Delta t} df$$

Since the discrete-time sequence was obtained by sampling the continuous-time signal, an important question is what is the nature of the relationship between G(•) and G_p(•). Now $g_t = g(t\Delta t)$ because we are sampling the continuous-time $g(t)$ every Δt seconds. We can represent g_t as $g_t = g(t\Delta t) = \int_{-\infty}^{\infty} G(f') e^{i2\pi f' t\Delta t} df$. This is equivalent to

$\sum_{k=-\infty}^{\infty} \int_{(2k-1)/2\Delta t}^{(2k+1)/2\Delta t} G(f') e^{i2\pi f' t\Delta t} df$. Now make a change of variable, setting $f' = f + k/\Delta t$. We can now rewrite the preceeding integral as $\sum_{k=-\infty}^{\infty} \int_{-1/2\Delta t}^{1/2\Delta t} G(f + k/\Delta t) e^{i2\pi (f + k/\Delta t) t\Delta t} df$. Now we bring the summation inside the integral to obtain. $\int_{-1/2\Delta t}^{1/2\Delta t} \sum_{k=-\infty}^{\infty} G(f + k/\Delta t) e^{i2\pi ft\Delta t} df$ by noting that $e^{i2\pi kt} = 1$ for all integers k, and t. So we have shown that

$$\int_{-1/2\Delta t}^{1/2\Delta t} G_p(f) e^{i2\pi ft\Delta t} df = \int_{-1/2\Delta t}^{1/2\Delta t} \sum_{k=-\infty}^{\infty} G(f + k/\Delta t) e^{i2\pi ft\Delta t} df$$

So $G_p(f) = \sum_{k=-\infty}^{\infty} G(f + k/\Delta t)$. The practical interpretation of this is that for any frequency f, $G_p(f)$ is the result of the contribution from an infinite set of frequencies $f \pm k/\Delta t$. In none mathematical terms, the amplitude (value) in the spectrum of the sampled process is is the sum of an infinite set of frequencies contained in the continuous-time process all

separated by integer multiples of the sampling frequency. This is referred to as aliasing. As a concrete example, consider a real-valued process that has a constant frequency spectrum from –1000 to 1000 Hz. Let us say that we have sampled this process at 2kHz and we wish to now investigate the amplitude of the process at 500 Hz (we are assuming we don't know that the spectrum is constant). From the above result, we know that the value of the discrete-time Fourier transform is the result of not only the value of the sampled process at 500 Hz, but also at $\sum_{k=-\infty}^{\infty} 500Hz \pm k2000$. Setting k=1 we see that 500 Hz receives contributions from 2500 Hz and –500 Hz. These are the contributions nearest to 500 Hz. However, we have stated that the continuous-time process contains no energy outside of ± 1000 Hz. Therefore the amplitude at 500 Hz in the discrete-time transform is equal to the continuous-time case (there are no additional contributions). Note that if we had sampled the process at only 1kHz, then 500 Hz would receive its nearest contributions from 1500 Hz and –500 Hz. In this case, the continuous-time process does contain energy at –500 Hz and that would be added to its value at 500 Hz to yield double the actual amplitude (because we have assumed that the value at –500 Hz and +500 Hz are equal). We would have accordingly created a spurious peak in the spectrum at 500 Hz as a result of aliasing. It can be shown in a rigorous manner that the if a continuous-time process is sampled at a rate at least twice as fast as the highest frequency present, then no aliasing will occur (this is the first scenario in our concrete example). This is called the *Shannon-Nyquist theorem*. In fact the highest frequency for which no aliasing occurs is $f = 1/2\Delta t$ and is referred to as the *Nyquist frequency*. This point is critical in any time and frequency-domain analysis of fMRI data because the typical sampling rates are not sufficient to

prevent respiration and cardiac frequencies to be unaliased. For example, consider an fMRI experiment in which the stimulus is presented at a period of 12 seconds (an unrealistically short period, but just for sake of argument). That period corresponds to a frequency of .0830 Hz. Further assume that our sampling interval is 3 seconds and that there are significant contributions in the data from the subject's breathing at .25 Hz (15 breaths/min). Since our sampling interval was 3 seconds we know that the value of the sampled spectrum at .0833 is equal to $.0833 \pm \sum_{k=-\infty}^{\infty} k.333$, because $1/\Delta t = .333$ in this case. Now since -.25+.333=.0830, we see that .0830 receives a contribution from the respiration frequency thereby creating energy at .0830 that could be mistaken for neuronal activation. In other words, even if we are sampling a pixel where there is no significant energy at our frequency of interest, .0830 Hz. The presence of significant energy at the respiration frequency would be aliased to .0830 Hz and could easily create a false positive.

Finally, because all practical applications of spectral estimation involve digital signals both in the time and frequency domains, it is important to discuss the case of discrete time and frequency. For example, we may be interested in the Fourier transform of a sequence $g_0, g_1, g_3, \ldots g_{N-1}$. For this sequence of length N we can define the Fourier transform to be another finite length sequence of numbers G_n, such that the G_n represent N equally-spaced frequencies. For simplification let us assume that the sampling interval Δt=1. The discrete Fourier transform of a sequence g_n is given by

$$G_k = \sum_{n=0}^{N-1} g_n e^{-i2\pi kn/N} \quad (13)$$

The inverse Fourier transform of G_k is given by

$$g_n = \frac{1}{N}\sum_{k=0}^{N-1} G_k e^{j2\pi kn/N} \quad (14)$$

Not only do these transform relations exists for finite-energy signals, but Parseval's relation, or Plancherel's theorem as it is sometimes called, guarantees that the energy in the time domain and frequency domain representations is equal. Formally, this is stated as follows in the continuos-parameter case. Analogous results exist for the discrete frequency/continous time, discrete time/continuous frequency, and discrete time/discrete frequency.

$$\int_{-\infty}^{\infty} |g(t)|^2 \, dt = \int_{-\infty}^{\infty} |G(f)|^2 \, df \qquad (15)$$

In theory we can represent a finite-energy signal exactly in time or frequency by means of the Fourier transform and this transform is energy-preserving. However, in time series analysis the actual form of the signal is not known and must be estimated from the observed data, just as the population mean is unknown and must be estimated in a more familiar problem. Furthermore, a stochastic process is typically not periodic, nor does it possess finite energy (because of the assumption of stationarity). Accordingly, it is not technically possible to represent a stochastic process by means of a Fourier series expansion or a Fourier integral. As a result, it is not immediately clear what it means to consider the spectrum of a random process. Fortunately, there is a result due to Cramer (1942), known as the spectral representation theorem, which allows the spectrum of a stationary process to be defined in such a way that its average power as a function of

frequency can be represented. Specifically, it can be shown that a wide-sense discrete-parameter stationary process may be represented as:

$$X_t = \int_{-1/2}^{1/2} e^{i2\pi ft} dZ(f) \quad (16)$$

where dZ(f) represents an orthogonal increment process (Loeve, 1963). The import of (11) is that a discrete-parameter stationary stochastic process may be represented as a weighted sum (integral) of sinusoids with random amplitudes |dZ(f)| and random phases-*arg*(dZ(f)). This result enables us to talk about the spectrum of a stationary process and accordingly about the average power in a process at a particular frequency, *f*.

3.2.1 The Cross-Correlation Method

As previously mentioned, time domain methods have been and continue to be widely used in fMRI data analysis (Bandettini et al., 1993; Cox et al., 1995). The low computational burden and apparent ease of interpretation of these methods are two chief reasons for their popularity. The most common time-domain method in fMRI involves computing the angle between a zero-mean vector or a series of zero-mean phase-shifted vectors and a zero-mean sampled signal. It is a simple matter to obtain a zero-mean vector from one with a non-zero mean by subtracting each component from the mean.

$$E\{x(i) - \bar{x}\} = E\{x(i)\} - \mu = \mu - \mu = 0 \quad (17)$$

Given two zero mean vectors we can compute the correlation coefficient by calculating the cosine of the angle between them. Let x and y be two zero-mean vectors, then the correlation between them is given by:

$$\frac{x^T y}{norm(x)norm(y)} \quad (18)$$

where the superscript T denotes vector transposition and the norm here is understood in the l^2 (finite discrete energy) sense, namely $\|x\| = \sqrt{x_1^2 + x_2^2}$ for a vector in R^2. This result is then easily generalized to R^n. It is easy to show that for zero-mean vectors the result in (18) corresponds to the more familiar measure of correlation, the Pearson product moment correlation coefficient. First note that the formulation in (18) is equal to the cosine of the angle between two vectors x and y. Secondly, recall that for any angle Θ the following equality holds: $-1 \leq \cos\Theta \leq 1$. The final step in the correspondence with the Pearson product moment correlation coefficient is to recognize that in (18) the mean is zero. For purposes of illustration the formula for the Pearson product-moment correlation coefficient is included here:

$$\rho = \frac{\sum_{i=1}^{N}(x_i - \mu_x)(y_i - \mu_y)}{\sqrt{\sum_{i=1}^{N}(x_i - \mu_x)^2}\sqrt{\sum_{i=1}^{N}(y_i - \mu_y)^2}} \quad (19)$$

Since we are interested in estimating the correlation coefficient from a sample, we will need to understand the sampling properties of the estimator, $\hat{\rho}$. Specifically, it will be shown the mean of $\hat{\rho}$ is zero and its variance given by 1/N.

Proof

Assume that $\{u_1, u_2, \ldots, u_n\}$ are normally-distributed with mean θ and variance σ^2. We will further assume the u_i's are identically and independently distributed (i.i.d) RVs. Further, assume we have a non-random vector given by $\{v_1, v_2, \ldots, v_n\}$. Furthermore, assume that both vectors are zero-mean. This can be done without loss of generality. We want the first and second moments of

$$\langle u,v\rangle / \|u\|\|v\|$$

Now, taking expectations

$$E[\langle u,v\rangle / \|u\|\|v\|]= \frac{1}{\|v\|}\sum_i E[\frac{u_i v_i}{\|u\|}] = \frac{1}{\|v\|}\sum_i v_i E[\frac{u_i}{\|u\|}]$$

since the u_i's are i.i.d the expected value of $\frac{u_i}{\|u\|}$ does not depend on *i*. Accordingly, we may choose any of the u_i since the expectation does not depend on *i*. So we choose u_1.

Finally, given our assumption of symmetry, we have the following:

$$E\left[\frac{u_i}{\|u\|}\right] = -E\left[\frac{u_i}{\|u\|}\right] = 0 \quad \text{for all } i.$$

Now we prove that the variance of this estimator is 1/N. Recall that the variance of a RV, *x* is given by $E\left[(x-\mu)^2\right]$. This is equivalent to $E\left[x^2\right]-\mu^2$. Note that we have assumed that our vectors are zero-mean, so we need only to compute the second moment given by:

$$E\left[\langle u,v\rangle^2 / \|u\|^2\|v\|^2\right] = \sum_{i,j} E\left[\frac{u_i v_i u_j v_j}{\|u\|^2\|v\|^2}\right] = \sum_{i,j}\frac{v_i v_j}{\|v\|^2} E\left[\frac{u_i u_j}{\|u\|^2}\right]$$

Again, assuming each u_i to be symmetric, we see that $E\left[\frac{u_i u_j}{\|u\|^2}\right] = 0$ if i≠j. Thus,

$$E\left[\langle u,v\rangle^2 / \|u\|^2\|v\|^2\right] = \sum_i \frac{v_i^2}{\|v\|^2} E\left[\frac{u_i^2}{\|u\|^2}\right]$$. Now given our assumption that the u_i's are i.i.d the expectation above will not depend on *i*. Then

$E\left[\frac{u_i^2}{\|u\|^2}\right]\sum_i \frac{v_i^2}{\|v\|^2} = E\left[\frac{u_i^2}{\|u\|^2}\right]$ since $\sum_i \frac{v_i^2}{\|v\|^2} = 1$. Now we compute

$E\left[\frac{u_i^2}{\|u\|^2}\right] = \frac{1}{n}\sum_i E\left[\frac{u_i^2}{\|u\|^2}\right] = \frac{1}{n}E\left[\frac{\sum_i u_i^2}{\|u\|^2}\right] = \frac{1}{n}$. This completes the proof. In summary we have seen that the mean of the estimator $\frac{\langle u,v\rangle}{\|u\|\|v\|}$ where one of the vectors is random, with a symmetric distribution and the terms are i.i.d is zero and its variance $\frac{1}{n}$.

3.2.2 Frequency-Domain Analysis

As previously mentioned, signals, which possess Fourier transforms may be characterized in terms of their frequency content. This involves projecting the signal onto a basis of complex exponentials with infinite support. In the case of stochastic processes this process is probabilistic since, as always, we are attempting to estimate population parameters from a finite sample. For the present, let us first consider a simple model:

$$X_t = \mu + \sum_f A(f)\cos 2\pi ft + B(f)\sin 2\pi ft \qquad (19)$$

Here the signal is represented as the sum of a constant term μ plus a combination of cosines and sines. Further, we assume that the amplitudes, A(f) and B(f), of each trigonometric function are uncorrelated RVs, such that $E[A_j,A_k]=E[A_j,B_j]=E[B_j,B_k]=0$. We also define $E[A_j^2]=E[B_j^2]=\sigma_j^2$. For such a process it can be shown that $E[X_t]=\mu$ and the process variance is simply the sum of the variances of the individual RVs (frequencies). If we define the spectrum as $S_j=\sigma_j^2$ for all j, then we can plot S_j vs. f_j and see where the

variability in the time series comes from. Viewed in this manner, the spectral analysis of a stationary process is really an analysis of variance technique. In practice S_j is referred to as the *power spectral density* (psd) or just *spectral density function* (sdf). Stated differently, a plot of the power spectral density function shows how the average power in a process is distributed by frequency. In statistical terms, the estimation of the spectral density function is concerned with the second moment nature of the process. Moreover, for each f_j the above model yields a two-degree-of-freedom estimate derived from two independent RVs, namely $A(f_j)$ and $B(f_j)$. Due to the limited number of degrees of freedom, it is reasonable to assume that any spectral estimate derived from such a model would exhibit considerable variability around the true spectrum. Obviously, if we were attempting to construct confidence intervals around our spectral estimate, it would be desirable to increase the degrees of freedom. It will be shown that this is nicely accomplished with the multitaper method.

3.2.3 The Autocovariance Sequence

It will be useful in the following discussion on spectral estimation to be armed with the notion of the autocovariance sequence (ACVS) of a time series. This definition will only be developed for the discrete case (hence the name sequence) but analogous results hold in continuous time. Recall that the covariance between two RVs X and Y are measured in the following fashion:

$$\sum_{t=1}^{N}(X_t-\mu_x)(Y_t-\mu_y) \quad (20)$$

Recall that (19) is an estimator for the corresponding population covariance

$$E[(X-\mu_x)(Y-\mu_y)] \quad (21)$$

Now assume instead of two RVs we have a time series of 100 values. Let $Y_t = X_{t+k}$ where k is an integer. If we adjust the summation in (19) to make use of all the data then the covariance between different points of the time series can be measured in the following way

$$S_r = \sum_{t=1}^{N-k} (X_{t+k} - \mu_k)(X_t - \mu_k) \quad (22)$$

Again, this is an estimate of the corresponding population ACVS given by

$$E[(X_t - \mu)(X_{t+k} - \mu)] \quad (23)$$

It will be helpful to illustrate the formulation in (21) with an example. Consider the sequence x_0, x_1, x_2, x_3. We could measure the covariance between x_0 and itself (designate it S(0,0)), x_0 and x_1, S(0,1) and so on. Note that the greatest lag between two elements in the sequence for which the covariance could be calculated would be 3, e.g. the covariance between x_0 and x_3, S(0,3). The above expression is referred to as the Autocovariance Sequence (ACVS) of a time series. Note that for S(0,3) the summation will run from 1 to N-3, so there will be only one product, namely (x_0x_3). If we were to construct a matrix of the autocovariance sequence in the above example, it would be the following

$$\begin{pmatrix} S(0,0) & S(0,1) & S(0,2) & S(0,3) \\ S(1,0) & S(1,1) & S(1,2) & S(1,3) \\ S(2,0) & S(2,1) & S(2,2) & S(2,3) \\ S(3,0) & S(3,1) & S(3,2) & S(3,3) \end{pmatrix}$$

Recall that if a process is stationary, the covariance between any two points depends only on the absolute difference between them. Accordingly, if the above sequence is stationary,

then the covariance S(0,1) and S(1,2) would be equal, as well as S(0,2) and S(1,3), etc. In fact the elements in the above matrix along any diagonal would be equal, with the elements of the main diagonal equal to the variance of the sequence. A matrix in which elements along any diagonal are equal is referred to as a *Toeplitz* matrix. It can be shown that a discrete-time random process is stationary if and only if its covariance matrix, S is Toeplitz. In addition to the Toeplitz property, it can be shown that the covariance matrix of a stationary stochastic process is *positive semi-definite*, or in most cases simply *positive definite*. A (LXL) matrix S is said to be positive semi-definite if for every (LX1) vector a the product $a^T Sa \geq 0$. The proof is as follows. Let y be the inner product of a and the observed vector x_n. Further use the Hermitian transpose operator H in taking the inner product so that

$$y=a^H x_n$$

Now, if a is (NX1) and x_n is (NX1) then y will be a scalar. Now note that

$$E[|y|^2]=E[yy^*]$$

$$=E[a^H x_n x^H a]=a^H E[x_n x^H]a=a^H Sa$$

where S is the autocovariance matrix for the sequence x_n. It will be shown below that the property of positive semi-definiteness plays an important role in determining how to estimate the ACVS of a finite sequence.

The importance of the ACVS in spectral estimation cannot be understated. In fact, in can be shown that the sdf of a random process is in fact the Fourier transform (discrete) of the ACVS, S_τ. (Stoica & Moses, 1997; Priestley, 1981).

3.2.4 Non-Parametric Spectral Estimation

Many common statistical problems involve the estimation of some unknown population parameters based on a finite sample. The same is true of the spectral estimation of a wide-sense stationary process. Namely, given a realization: x_1, x_2, x_3, ... x_n of a wide-sense stationary process we are interested in how the autocovariance sequence (ACVS) can be estimated. Recall that the sdf is given by the discrete-time Fourier transform of the ACVS:

$$S(f) = \Delta t \sum_{\tau=-\infty}^{\infty} s_\tau e^{-i2\pi f \tau \Delta t} \quad (24)$$

Furthermore, the ACVS is given by

$$E[(X_t - \mu)(X_{t+\tau} - \mu)] \quad (25)$$

From the above expression, it is clear that the population parameter μ must first be estimated. As in more familiar estimation problems, this quantity can be estimated as:

$$\overline{X} = \frac{1}{N}\sum_{t=1}^{N} X_t \quad (26)$$

With the linearity of the expectation operator, it can be shown that this estimate is unbiased, i.e.

$$E[X_t] = \frac{1}{N}\sum_{t=1}^{N} E[X_t] = \mu \quad (27)$$

Given a zero-mean, wide-sense stationary process, the ACVS may be estimated in the following manner

$$\hat{S}_\tau = \frac{1}{N-|\tau|}\sum_{t=1}^{N-|\tau|}(X_t - \overline{X})(X_{t+\tau} - \overline{X}) \quad (28)$$

If μ is substituted for the sample mean in the above expression, it can be shown that (28) is an unbiased estimator of the ACVS. Another estimator of the ACVS of a zero-mean, wide-sense stationary process is given by

$$\overset{(b)}{S}_{\tau} = \frac{1}{N}\sum_{t=1}^{N-|\tau|}(X_t - \overline{X})(X_{t+\tau} - \overline{X}) \quad (29)$$

The superscript (b) here designates that this is a biased estimate of the ACVS. If μ is substituted for $\overline{X}$ in the above expression and the expectation operator is applied, the result is equal to

$(1 - \frac{|\tau|}{N})s_{\tau}$. Therefore, this estimator is biased. In practice, the biased estimator is preferred for the following reasons:

1.) It can be shown the mean squared error (MSE) expressed as MSE=variance +(bias)^2 is smaller for the biased estimator. In practical terms this says that the variance in the unbiased estimate is more harmful than the bias in biased estimator. This effect is particularly marked for $|\tau|$ close to N. For example, consider the two estimators for τ=N-1. In this case, the expression for the unbiased estimator becomes

$$(X_1 - \overline{X})(X_N - \overline{X}) \quad (30)$$

The analogous expression for the biased estimator yields

$$\frac{1}{N}(X_1 - \overline{X})(X_N - \overline{X}) \quad (31)$$

It is clear from the above that the variance of the unbiased estimator for τ=N-1 is N^2 times that of the biased estimator.

2.) It can be shown in the theory of stationary processes that the ACVS must always be positive semidefinite (Percival & Walden, 1993). While the sequence obtained with the

biased estimator is always positive semidefinite, the sequence for the unbiased estimate need not be.

3.) It can be shown for a process, X_t with a purely continous spectrum that the true ACVS decreases to 0 as $|\tau|$ approaches N-1. As can be seen in (30-31) above, the biased estimator is closer to 0 at N-1 than the unbiased estimator.

3.2.5 The Periodogram

Let us assume that we have a discrete-parameter zero-mean real-valued stationary process. Given a realization of this process: $x_1, x_2, x_3 \ldots x_n$, it has been shown that the ACVS can be estimated (in a biased case) by

$$\frac{1}{N}\sum_{t=1}^{N-|\tau|} X_t X_{t+\tau} \qquad (32)$$

If we now substitute the above for s_τ and use the finite sum we can obtain an estimator of the sdf.

$$\hat{S}(f) = \Delta t \sum_{\tau=-(N-1)}^{N-1} \hat{s}^{(b)}_{\tau} e^{-i2\pi f \tau \Delta t} \qquad (33)$$

Through substitution and a change of variable, the following expression can be derived

$$\hat{S}(f) = \frac{\Delta t}{N}\left|\sum_{t=1}^{N} X_t e^{-i2\pi f t \Delta t}\right|^2 \qquad (34)$$

This expression for the sdf of a stationary process is typically referred to in the literature as the *periodogram*. Now, assume that $\Delta t = 1$. Define $J(f) = \frac{1}{\sqrt{N}}\sum_{t=1}^{N} X_t e^{-i2\pi f t}$ for $|f| \le \frac{1}{2}$.

Note that (34) can be written as $|J(f)|^2$. By using the spectral representation theorem, it can be shown that

$$E[|J(f)|^2 = N \int_{-1/2}^{1/2} \int_{-1/2}^{1/2} e^{-i(N+1)\pi(f''-f)} \frac{\sin(N\pi(f''-f)}{N\sin(\pi(f''-f)} e^{i(N+1)\pi(f'-f)} \frac{\sin(N\pi(f'-f)}{\sin(\pi(f'-f)} E[dZ^*(f'')dZ(f')]$$

where dZ(f) represents an orthogonal increment process The import of the above representation is that the use of the periodogram as an estimate of the sdf leads to a convolution of the true sdf with $\frac{(\sin(N\pi f))^2}{N(\sin(\pi f))^2}$. Let us denote this expression as $F_n(\bullet)$ and refer to it as *Fejer's kernel*.

If we plot Fejer's kernel for various choices of N, we will see the following:

- For any choice of N, Fejer's kernel will have a broad central peak which is symmetric about the origin and peaks in magnitude at f=0.
- In addition to the peak at f=0, Fejer's kernel will have N-2 sidelobes which decrease in magnitude as $|f|$ increases. The magnitude of the first sidelobe is about 13 dB below that of the central lobe.
- As N goes to infinity, the N in the denominator dominates and Fejer's kernel will approach a Dirac delta function at f=0. Another way to state this is that it will become asymptotically unbiased.

In more descriptive terms, the convolution of the sdf with Fejer's kernel means that the spectrum of our random process has been contaminated by the spectrum of the kernel, which is entirely an artifact of the computational procedure. The result of this contamination is a loss of spectral resolution and contamination of the spectral estimate by the sidelobes of the kernel. The loss of resolution is a result of spectral "smearing" caused by the width of the main lobe in the kernel's frequency –domain representation. If two sinusoidal components in the underlying sdf are closely spaced, they may be smeared

together due to the broad frequency profile of the main lobe. The second source of concern comes from a transfer of energy from the sidelobes. In other words, the spectral estimate may contain significant energy at points simply due to the sidelobes of the kernel. This effect on the spectral estimate is referred to in the literature as *leakage*. The combined effect of a loss of resolution and leakage is that one cannot be certain if spectral components in the estimate are a true part of the spectrum or simply an artifact of the convolution with the kernel.

Because we have only a finite number of terms from the underlying process we are only able to examine the process through a narrow *window*. The more data points we have, the wider the window, but still the process is inevitably truncated abruptly at some pair of points say (–M) and M. The frequency domain representation of this rectangle is Fejer's kernel. From the discussion above, it is clear that the "ideal" windowing function in the frequency domain would have a narrow main lobe and minimal side lobes. One way to alleviate some of the problems encountered with Fejer's kernel is to introduce a different windowing function with smoother transitions. Since a jump discontinuity in a function introduces ripples in the Fourier transform, if we are able to create a windowing function with no discontinuities, i.e. differentiable, than the resulting sidelobes would be reduced and the derived spectral estimate may not be as badly biased. Windowing a time series amounts to multiplying it point by point by the time-domain representation of another function. This process is referred to as *tapering*. While tapering reduces the size of the sidelobes encountered in the convolution with the Fourier transform of the windowing function, it engenders a number of other serious problems. When you taper a data set, you effectively lose data. This is because you are weighting values in the time series

differentially as a result of the multiplication with the windowing function's time domain representation. It is a general principle of statistical estimation that the weight given a particular value should only decrease or increase with a change in variance, i.e. the more variable some quantity is, the less we are certain of its true value, and accordingly, the less weight it should be given. Yet, the assumption with time series data has been that the variance is independent of time, i.e. the same for all points. Still, several windowing functions have been found with sidelobe properties superior to those of the rectangular window. Table one lists some common windowing functions with their time-domain generating equations, the width of the main lobe, and the magnitude of their most prominent sidelobe relative to the main lobe in dB. All the windowing functions w(k) satisfy the following properties: w(k)=0 for |k|>M, where M is some integer, and w(k)=w(-k), i.e. they are even functions.

Table 1-Common windowing functions

Window Name	Generating Equation	Main Lobe width in Radians (approx)	Sidelobe level in dB
Rectangular	$w(k)=1$	$2\pi/M$	-13
Bartlett	$w(k)=(M-\|k\|)/M$	$4\pi/M$	-25
Hanning	$w(k)=.5+.5\cos(\pi k/M)$	$4\pi/M$	-31
Hamming	$w(k)=.54+.46\cos(\pi k/(M-))$	$4\pi/M$	-41
Blackmann	$w(k)=.42+.5\cos(\pi k/(M-1))$ $+.08\cos(2\pi k/(M-1))$	$6\pi/M$	-57

In addition to tapering a time series to produce more desirable leakage properties, it is possible to reduce the variance of a spectral estimate by averaging over multiple windows or multiple segments. This process is analogous to reducing variance by increasing sample size in estimating a population mean. Several multiple window techniques have been developed for non-parametric spectral estimation (Percival & Walden, 1993; Thomson, 1994). One of the most popular is *Welch's overlapped segment averaging* (WOSA). In this approach, the data record is divided into overlapping segments, a data taper is applied to each segment (typically a Hamming or Hanning window) and the Fourier transform is computed. These Fourier transforms are then averaged to produce the spectral estimate. In multitaper spectral analysis (Thomson, 1982), the data record is not segmented. Rather, several different data tapers are used and the Fourier transform is computed for the whole data set with each taper. Finally, these spectral estimates are then averaged together.

An important point to consider in multiple window techniques is that the degree of variance reduction crucially depends on how correlated the windowing functions, or segments are. To illustrate this, consider attempting to estimate the mean height of all male students at the university. A typical approach would be to randomly sample male students and measure their heights. Let (Y_1, Y_2, Y_3, ... Y_n) denote our samples. Now if the samples are independent and taken from a normal distribution with mean, μ and variance, σ^2, it can be shown that $\overline{Y}$ is normally distributed with mean

$$E\{\overline{Y}\} = E\{\frac{1}{n}\sum_{i=1}^{n} Y_i\} = \mu \quad (35)$$

and variance

$$V\{\overline{Y}\} = V\{\frac{1}{n}\sum_{i=1}^{n} Y_i\} = \frac{\sigma^2}{n} \quad (36)$$

The important result from the above is that if independent samples are taken, then the variance of an estimate will decrease as the sample size increases. A similar situation holds in producing multiple window spectral estimates. Specifically, if the windows (or segments) are uncorrelated, then averaging over multiple windows will decrease the variance. It will be seen that this requirement is fulfilled if the windowing functions used in multitaper spectral estimation are *discrete prolate spheroidal sequences*.

CHAPTER 4
MULTITAPER SPECTRAL ESTIMATION

4.1 Slepian Sequences

It is evident from table one that the sidelobes are dramatically reduced in non-rectangular windowing functions. Recall that this is desirable because it produces less leakage in the resulting spectral estimate. However, note that while the sidelobes decrease the width of the main lobe increases. As previously mentioned, this results in reduced resolution, i.e. the ability to resolve closely spaced spectral components. On the basis of that consideration alone, it might be tempting to ask whether or not windowing functions can be found with more optimal mainlobe-sidelobe properties. At the heart of this problem lies the fact that time and frequency are not independent. Recall that the time and frequency representations of a signal are linked via the Fourier transform. A well know result known as the scaling property (see Bracewell, 1988) states that if a function $\sqrt{a}g(at)$ possesses Fourier transform, its transform will be given by $\frac{1}{\sqrt{a}}G(\frac{f}{a})$. This result is easily proved. The Fourier transform of $\sqrt{a}g(at)$ is defined to be

$$\int_{-\infty}^{\infty}\sqrt{a}g(at)e^{-i2\pi ft}dt \qquad (37)$$

Now do a change of variable by setting $y = at$, then $dy = adt$. Making the proper substitution in (37), we obtain

$$\frac{1}{\sqrt{a}}\int_{-\infty}^{\infty} g(y)e^{-i2\pi(f/a)y}dy = \frac{1}{\sqrt{a}}G(\frac{f}{a}) \qquad (38)$$

In non-mathematical terms the interpretation is that if one member of the transform pair contracts horizontally and expands vertically, the other member expands vertically and contracts horizontally. An important implication which can derived from (37-38) and other results in Fourier analysis is that a function cannot be simultaneously band-limited (frequency domain) and time-limited. The more we squeeze a function or stretch it in one domain the more spread out or concentrated it will become in the other. To readers familiar with quantum mechanics, this is an example of Heisenberg's uncertainity principle. In real world problems we are inevitably. dealing with time-limited (or in the discrete case "index-limited") functions since we cannot handle an infinite number of samples. This means that our functions will not be band-limited. Fortunately, the Fourier transform of a finite number of samples of a continuous-time real, or complex-valued signal is periodic. Otherwise, we would not be able to represent an index-limited sequence in the frequency domain, because its spectrum would be infinitely wide. Also, if we are dealing with band-limited functions, we will not be able to find one, which is index-limited. However, it is possible to ask which elements of the set of functions band-limted to [-w,w] have most of their energy concentrated in some index set [0,N-1]. In other words, it is understood that a function cannot be found which is band-limited to [-w,w] and simultaneously index-limited to [0, N-1], but we can attempt to solve the following maximization problem: Assume we have a class of functions band-limited to [-w,w] such that

$$\beta^2(w) = \int_{-w}^{w} |G_r(f)|^2 df \Big/ \int_{-\infty}^{\infty} |G_r(f)|^2 df = 1 \qquad (39)$$

The above expression illustrates that the total energy in the band [-w,w] is the same as in the band [-∞, ∞]. Now consider the following ratio in the time domain:

$$\alpha^2(T) = \sum_{t=0}^{N-1} |g_t|^2 \Big/ \sum_{t=-\infty}^{\infty} |g_t|^2 \qquad (40)$$

If the g_t in (35) are considered to be the Fourier transforms of the $G_p(f)$ in (39), then the ratio in (40) must be less than unity. In other words, all the energy cannot be concentrated in the index set [0,N-1]. Still, it is possible to find g_t such that the ratio in (40) is maximized. Conceptually, we are asking for the class of functions band-limited to [-w,w] which ones have most of their energy concentrated in the time-domain in [0, N-1]. This problem was addressed in a series of seminal papers for both the continuous and discrete cases by Slepian (1976, 1978, 1983). It can be shown that in the discrete case (see Slepian, 1978 these optimally time and band-limited functions are eigenfunctions of the following integral equation

$$\int_{-w}^{w} \frac{\sin(N\pi(f'-f))}{\sin(\pi(f'-f))} H_r(f)df = \alpha^2(N) H_r(f') \quad |f'| \le w \qquad (41)$$

The above is an example of a Fredholm integral equation with a degenerate kernel. It follows that this equation will have a finite number (N) of eigenvalues and accompanying eigenfunctions. Furthermore, these eigenvalues will be real-valued, positive and all less than one. They may also be ordered such that

$$1>\lambda_0(N,W)>\lambda_1(N,W)>\lambda_2(N,W)>\ldots>\lambda_{N-1}>0 \qquad (42)$$

The eigenvalue $\lambda_0(N,W)$ and its eigenfunction will provide the maximum value to the earlier concentration problem. Similarly, the eigenvalue $\lambda_1(N,W)$ and its corresponding eigenfunction will provide the next highest concentration of energy in (0,N-1) and the

eigenfunction corresponding to λ_1 will be orthogonal (uncorrelated) with that eigenfunction corresponding to λ_0, etc. Continuing in this manner it can be shown that these N eigenfunctions form a complete orthogonal set, which can be easily normalized to yield a complete orthonormal set. The eigenfunction corresponding to λ_k is referred to as the kth-order discrete prolate spheroidal wavefunction (DPSWF), while its time-domain counterpart is referred to as the kth-order discrete prolate spheroidal sequence (DPSS).

The significance of these functions in spectral analysis has been primarily in their utility as optimal windowing and averaging functions because of the following two properties:

1.) The DPSSs are orthogonal in time and their Fourier transforms are orthogonal in frequency. From this it follows that they are uncorrelated (or nearly so) and as such provide a good basis for a multiple-window technique.

The DPSS have excellent mainlobe-sidelobe properties, which make them ideal windowing functions. The sidelobes of the DPSSs for k=0 and 1 are -100 dB and -80 dB respectively.

4.2 The Multitaper Spectral Estimator

As a way of alleviating the problems encountered in tapering and simultaneously reduce variance, Thomson (1982) introduced an approach that involves the use of multiple tapers, which are orthogonal both in the time and frequency domains. This has become referred to in the literature as multitaper spectral estimation (Percival & Walden, 1993). Since its development, multitaper spectral estimation has been applied to time series data in geophysics (Park et al., 1987), the study of solar rotation frequency (Denison et al., 1998), and magnetic field data (Thomson et al., 1995). In its simplest formulation, the multitaper spectral estimate is the average of K direct spectral estimators

$$S_{mt}(f) = \frac{1}{K}\left|\sum_{k=0}^{K-1} h_{k,N} X e^{-i2\pi fk}\right|^2 \quad (43)$$

Orthogonality ensures that the spectra of the tapers are uncorrelated, or at least very nearly so. Computing K Fourier transforms, or *eigenspectra*, over the K uncorrelated data tapers with equal variance produces an estimate of the spectrum that has less variance than any of the tapers considered alone. It has also been shown in a direct comparison with WOSA, that multitaper spectral estimation reduces the variance of the estimate by 50-75% and leakage by 10-20 dB (Bronez, 1992). Additonally, the multitaper method leads to other advantages as well, including an appealing statistical test for the presence of sinusoidal components in noise. It is this application that we are most interested in for the analysis of fMRI data.

The utility of the multitaper spectral estimator for fMRI data results from the following:

1.) The multitaper method provides a non-subjective way to set the resolution of a spectral estimator. This is rather difficult to do for spectral estimators in general (Kay & Demeure, 1984).

2.) The multitaper method allows for generation of spectral estimates with multiple degrees of freedom. This in turn shrinks the size of corresponding confidence intervals considerably.

3.) There exists an appealing statistical test for the presence of sinusoidal (line) components in additive noise based on complex regression theory.

The simplest recipe for computing the multitaper spectral estimate requires the following steps:

1.) Select the resolution bandwidth.

2.) Compute the eigenspectra (43) for the Slepian sequences k=0,1…K-1 with K dependent on the resolution bandwidth.

3.) Average together the K eigenspectra

Points 1 & 2 above require some explanation. The resolution bandwidth dictates how accurately closely spaced spectral peaks will be identified. The resolution of a discrete time series is basically inversely proportional to the sample size. For example, with n=128 the best possible resolution is $\frac{1}{128}$. This presents a potential problem in fMRI because sample sizes are typically small, e.g. ≤100 points. Therefore, in the case of n=100, line components less than 0.01 cycles/sample apart will not be well resolved. In the multitaper method, the first step is to choose the resolution, W so that W is a small multiple of the fundamental frequency, $W = \frac{j}{N\Delta t}$ where N=sample size, and Δt is the sampling period and j is a positive integer. Here j is typically taken to be some small number, e.g. 2,3, or 4. In the case of a 100-point time series assuming Δt=1, the resolution becomes $\frac{2}{100} = 0.02$. This immediately begs the question, why not choose $\frac{1}{100}$ to make the resolution as good as possible. The answer is that in the multitaper method the number of DPSSs with good leakage properties depends on W. The number of tapers, K should be chosen so that K<2NWΔt (Percival & Walden, 1993). For the case of j=2 above, this would result in at least the first three Slepian sequences providing good leakage properties, i.e. K=0,1, and 2. If we take j=1. Then only the zero-th order Slepian sequence will exhibit good leakage properties. In this case the multitaper method will reduce to the computation of a direct spectral estimate with the zero-th order DPSS as the windowing function and the averaging properties will be lost.

4.3 Line Detection in the Multitaper Method

A reasonable first assumption about fMRI data is that it can be modeled as a sinusoid or sum of sinusoids in a colored noise background, i.e the spectral density function of the background may not be white noise (flat spectrum). The multitaper method provides a convenient test for such signals based on complex regression theory. The resulting test statistic as derived in Thomson(1982) is analogous to the usual F-test used in real-valued regression models. It will be useful to briefly review the derivation of this statistic in a more familiar setting and then extend it to the detection of a sinusoid in noise.

Recall that in linear regression you are attempting to fit a model of the form

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots \beta_i x_i + \varepsilon_i \quad (44)$$

In matrix notation the above can be re-written

$$Y = X\beta + \varepsilon \quad (45)$$

where X is the independent variable matrix, β is the matrix containing the model parameters and ε is the matrix of the error terms ε_i, or departures from a deterministic linear model.

This is referred to as a linear model because the expression is linear in β. The ε-term is typically held to be a Gaussian-distributed RV with zero mean and constant variance. Because the model contains a RV, we cannot predict exactly what the value of the dependent variable will be for any independent variable. However, by using the theory of least squares the best estimates of β_i can be computed and the fit of the corresponding model tested. Specifically, we are interested in testing the null hypothesis that all the β_i are

equal to zero vs. the alternative hypothesis that at least one β_i is not. The test statistic used for this is essentially the following ratio:

F value=explained variability(accounted for by model)/unexplained variability

If the numerator in the above is large enough, then it will be accepted that the proposed model captures something about the "real" relationship between the independent and dependent variables. In an attempt to place this on more of a mathematical foundation the above can be stated (yet still verbally) in the following manner:

$$\text{F value}=\frac{\text{Sum of squares(model)/df(model)}}{\text{Sum of squares(error)/df(error)}}$$

where df denotes degrees of freedom

Since the sum of squared Gaussian RVs is χ^2-distributed and the ratio of two chi-square RVs is F-distributed, the F-distribution may be used to determine the proper confidence intervals for the above test statistic (Freund, 1992)

In many signal detection applications, a similar model to the above is formulated, namely we assume that the value at any point is a combination of a sinusoidal component(s) plus an error term (noise). Specifically, in the following discussion it will be assumed that the amplitude, frequency , and phase angles of the sinusoidal components are unknown constants, i.e. they are not RVs. It will be further assumed that the noise has a Gaussian distribution, but that its frequency spectrum is not constant, i.e. it is colored. Now, consider the simplest formulation of this model, the signal given by:

$$X_t= D_1 \cos (2\pi f_1 t + \phi_1) + \eta_t \quad (46)$$

Where D_1, f_1, and ϕ_1 are unknown constants and η_t is a stationary zero mean continuous background of noise with a sdf $S\eta$ (•) which is not necessarily constant. Then applying the expectation operator the following expression is obtained:

$$E\{X_t\}=E\{ D_1 \cos (2\pi f_1 t + \phi_1) + \eta_t\}= C_1 e^{i2\pi f1t} + C_1^* e^{-i2\pi f1t} \quad (47)$$

where $C_1=(D_1/2)e^{i\phi 1}$ and * represents complex conjugation. Now consider the following expression:

$$J_k(f)=\sum_{t=1}^{N} h_{t,k} X_t e^{-i2\pi ft} \quad (48)$$

In the above expression $h_{t,k}$ represents the k^{th} Slepian sequence (data taper) and $J_k(f)$ represents the Fourier transform of the data sequence multiplied by the kth Slepian sequence. Now let $H_k(f)$ represent the Fourier transform of $h_{t,k}$ given by $H_k(f)=\sum_{t=1}^{N} h_{t,k} e^{-i2\pi ft}$. It can be shown that $E\{J_k(f)\}=C_1H_k(f-f_1) + C_1^*H_k(f+f_1)$. Also note that $J_k(f)$ can be written in terms of its expected value as $J_K(f)=E\{J_k(f)\} + \sum_{t=1}^{N} h_{t,k}\eta e^{-i2\pi ft}$. By setting $f=f_1$ we find that $E\{J_k(f_1)\}=C_1H_k(0) + C_1^*H_k(2f_1)$. For a function concentrated in a particular band, say [-w, w], if $2f_1>w$ then $H_k(2f_1) \approx 0$ and $E\{J_k(f_1)\}=C_1H_k(0)=C_1\sum_{t=1}^{N} h_{t,k}$. With the assumption that $2f_1>w$ and that $\varepsilon_k=\sum_{t=1}^{N} h_{t,k}\eta e^{-i2\pi ft}$ has a complex Gaussian distribution, then $J_k(f_1)=C_1H_k(0) + \varepsilon_k$ defines a complex-valued regression model. The problem of detecting a sinusoid in a continous background is now formulated as testing the null hypothesis that $C_1=0$ vs. the alternative that $C_1\neq 0$. Note that under the null hypothesis $J_k(f_1)=\varepsilon_k$.

An estimator for C_1 can be found using the theory of complex least squares. This amounts to finding $\hat{C}_1$ that minimizes

$$\sum_{k=0}^{K-1} \left| J_k(f_1) - \hat{C} \frac{H_k(0)}{(\Delta t)^{1/2}} \right|^2 \quad (49)$$

The F variance-ratio test (after some reduction) is given by

$$\frac{(K-1)|C_1|^2 \sum_{k=0}^{K-1} H_k(0)}{\Delta t \sum_{k=}^{K-1} \left| J_k(f_1) - \hat{J}_k(f_1) \right|^2}$$

where $\hat{J}_k(f_1) = \frac{\hat{C}_1 H_k(0)}{(\Delta t)^{1/2}}$. The above statistic is F-distributed (the ratio of two χ^2 RVs) with 2 and 2K-2 degrees of freedom. If $C_1 \neq 0$ then then the value of the test statistic will exceed some significant value on that F-distribution, say 99%. In signal terms the F-ratio above depicts the ratio of total power explained by the assumption of a line component at a given frequency to the residual power. These results are naturally extended to the detection of multiple lines via a partial F-test (Draper & Smith, 1998).

An important point to make here is that when testing whether the amplitude of each frequency is different from zero, we are forced to make N tests, where N is the number of frequencies. This requires us to be careful in our specification of the level, α, at which we will reject the null hypothesis. Specifically we are interested in the probability that at least one of the intervals we construct will fail to contain the true amplitude for any given frequency, f. This is given by

P[at least one of the N intervals fails to contain true amplitude]=1-P[all N intervals contain the true amplitude]$\geq$ 1-(1-α)$^N \geq \alpha$. Therefore in order to make the probability of at least one false detection (significant non-zero amplitude) equal to α, we must specify the individual levels of significance to be less than α. There are many methods for determining this level of significance. The simplest formulation is the Bonferroni procedure (McClave & Dietrich, 1994). The Bonferroni procedure calls for specifying the significance level for each comparison at α/N, where N is the number of comparisons. This is the method to be adopted here.

CHAPTER 5
APPLICATION TO SIMULATED AND ACTUAL FMRI DATA

Armed with some background for data analysis based on the cross-correlation (XCR) method and the multitaper spectral estimator (MTM), we will now examine the applications of these methods to actual and simulated fMRI data. The following section describes the conditions under which the simulated fMRI data was collected and analyzed.

5.1 Application to Simulated Data

Monte Carlo simulations were generated and analyzed in MATLAB with programs written by the author. The time-domain cross correlation was run using the routine described in Cox & Jesmanowicz, 1995. Five reference vectors were used in each XCR analysis. The five reference vectors were lagged in each case so that they spanned ¼ period ($\pi/2$ radians) of the forcing frequency. Multitaper spectral analysis was carried out on the time series data using algorithms described and employed in Mann, 1997 and Mann & Park, 1994. Several parameters were varied during the simulations including the signal-to-noise ratio (SNR), the spectra of the additive noise, and the sample size. The probability of a type I error was specified for the respective methods in the following manner. In the MTM method, the significance level for the detection of a non-zero amplitude was set at (1-1/N). This was chosen in accordance with the literature on multitaper spectral estimation and multiple regression (Draper & Smith, 1998; Thomson, 1982). The chief purpose of this specification is to guard against the possibility of at least

one false detection in N comparisons in accordance with Bonferroni procedure discussed at the end of the last chapter. Significant correlations in the time domain were set to yield a probability of 1.0×10^{-6}. The latter threshold was chosen because it corresponds to a correlation of $|0.5|$ with a sample size of 100 points. The majority of the simulations to be discussed below were carried out with a sample size of 100, a number representative of series size in actual fMRI experiments. It should be kept in mind that this number is completely arbitrary and the sample size could just as easily been 68, or 80 for example. Both methods assume that the signal may be modeled as a sinusoidal component (line component) in additive noise. The cross-correlation method further assumes that this line component may be phase-shifted by some amount τ. The simulated data reported consist of the following:

1. Matrices of 64X64X100 were generated consisting of single line component added at pixel [38,39]. White Gaussian noise was added to all pixels. One simulation was performed with AR(1) (red) noise only (no added signal). The AR(1) noise is the result of convolving Gaussian noise with an all-pole (IIR) filter. The AR(1) model used has been suggested as a realistic model for the noise encountered in fMRI (Purdon & Weisskoff, 1998).The variance of the white and red noise were generated to be consistent with observed samples of actual fMRI data.

2. The general signal conditions including: frequency of interest, number of samples, and sampling rate are consistent with the data acquisition in an auditory fMRI study conducted by the author (Mohr et al, 1999).

3. All images were written to disk so that the exact noise conditions were used with both the cross-correlation and multitaper method.

4. Phase shifts used in the onset of the line component are in agreement with published data on hemodynamic delays (Bandettini et al., 1993).

5. Simulations examining the effect of sample size were performed by maintaining the identical sampling rate and SNR conditions, and altering only the number of samples, i.e. (64X64X550, 64X64X450,

etc.). These conditions were identical to the ones utilized in the collection of the experimental data.

The first two simulations were run using white noise added to a single line component-sin($2\pi 0.05t$). The amplitude of the sinusoidal component in the first simulation was fixed at 4% above the amplitude of the noise. Furthermore the sinusoid was phase-shifted to mimic a hemodynamic delay of two samples (6 seconds) in the first simulation and one sample (3 seconds) in the second. Five reference boxcar vectors were created with periods equal to the sinusoid (20 samples) for the time-domain correlations.

Fifty simulations were run with a single line component at 0.05 cycles/sample and with an amplitude 4% above the noise floor and phase-lagged two samples. In 38 of the 50 (76%) simulations the cross-correlation method detected the presence of the line component at $|0.5|$ or above. In the remaining 12 trials the cross-correlation method failed to detect the sinusoid. The multitaper spectral estimator detected line components at 0.05 in 46 of the 50 (92%) trials.

A trial by trial comparison with cross-correlation method indicates that both detection methods missed the line component in trials 4 and 27, while the cross-correlation method detected line components in trials 10 and 26 where the MTM estimator did not. There was detection agreement between the methods in 38 or the 50 trials (76%).

The next simulation was run using a phase lag of one sample instead of two. This corresponds to a hemodynamic delay of 3 seconds under the given sampling scheme. With the exception of this phase difference, the conditions were identical to those in simulation one. In the second simulation the cross-correlation method again detected the line component in 38 of 50 trials (76%). The MTM estimator detected the line component

in simulation II in 45 of 50 (90%) trials. In this simulation the two methods agreed on 35 of 50 trials (70%). In both of the simulations with a single line component 4% above the noise floor, the MTM was more accurate in detecting the line component than the cross-correlation method. Specifically, the MTM estimator failed to detect the line component in 9/100 (9%) trials, while the cross-correlation rejected the presence of a line component in 24/100 (24%) trials.

The fact that the overall detection performance between estimators (percentage correct) did not change is not surprising given that the only difference between simulation I and II was a one-sample phase lag. The multitaper method employed here makes no use of phase information and therefore should not be affected by any phase lag. The XCR method utilized five reference vectors, which could cover a phase lag of up to5 samples (15 seconds, or ¼ of a cycle of the signal vector). Accordingly, a phase-lag of one or 2 samples would not be expected to make a significant difference.

While the first two simulations used white noise as the background, the appropriateness of this model for fMRI data has been questioned by a number of writers (Mitra, 1997; Purdon & Weisskopf, 1998). A recent study has suggested that the noise in fMRI data is best modeled as red noise, an autoregressive process of order one –AR(1). An autoregressive process of order 1 is defined by the difference equation:

$$X_t - aX_{t-1} = \varepsilon \quad (50)$$

where ε_t is a Gaussian process.

It can be shown that if $|a| < 1$ an AR(1) process will be stationary. In the frequency domain, an AR(1) process has most of its energy in the low frequencies, this is in contrast

to a white noise process where the spectrum is flat (equal power at all frequencies. Employing parameters derived by Purdon & Weisskopf, 1998, a 100-trial simulation of 64X64X100 was conducted with AR(1) noise only, i.e. no line components present. Again, signal intensity and variance were constructed to be consistent with observed fMRI data. Specifically, the gain and coefficients of the AR(1) process were obtained by fitting an AR(1) model to fMRI data obtained using EPI at 3-Tesla on the same Signa scanner used in the actual experiment. Again, the random matrices were written to disk and subjected to analysis by both the cross-correlation and MTM methods. In 100 trials the cross-correlation method found a total of 13 pixels with a correlation of $|.5|$ or higher. The MTM method found 2 pixels in the 100-trials with a significant frequency component at .05. Overall, there were 18 pixels found by the MTM method with at least one significant line component over all the possible frequencies.

It is possible that a line component may be present at even poorer signal-to-noise ratios than the 4% criterion used in simulations I, and II. For example, it has been reported that because of the effect of gradient noise during data acquisition, the dynamic range of auditory activation is reduced (Talavage et al, 1999). Therefore, it is critical that data-analysis methods be effective at a variety of SNRs. To address this point, a simulation was run repeating the conditions in simulation II except that the amplitude of the line component was reduced to 2% above the amplitude of the white noise. The cross-correlation method failed to detect the line component at the |0.5| level in 10/50 trials (20%). This corresponds to a correct-detection rate of 80%. While the cross-correlation method failed to detect the sinusoid in 10/50 trials, the MTM estimator correctly detected in the sinusoid in 46/50 (92%) of the trials.

Continuing to examine the effect of degrading the SNR, the final 100-sample point simulation utilized a sinusoid with an amplitude 2% below the mean value of the noise. The cross-correlation method failed to detect the line component in 14 of the 50 (28%) trials. The MTM spectral estimator was able to detect the line component even though its amplitude was 2% below the noise in 45/50 (90%) of the trials. Alternatively, it failed to detect the line component in 10% of the trials. Contrast this with the performance of the time-domain method in which its detection performance was 36/50 (72%).

In the above simulations, the sampling rate and number of samples were held fixed while the two parameters of the line component, amplitude and frequency were manipulated. The next simulations address what the effect of fewer periods of the signal (fewer samples) has on data analysis.

In order to assess the effect of sample size on the signal detection, matrices of 64X64X550 were generated. Again, a line component was inserted at pixel [38,39]. In this instance the frequency of the line component was .0167 cycles/sample. The matrices were written to disk and subsequently data points were deleted from the end of the images, creating five new datasets: 64X64X500, 64X64X450, 64X64X350, 64X64X250, and 64X64X150. With the stated frequency this consists of the following number of periods: 1.) 550 points≈9, 2.) 500≈8.5, 3.) 450=7.5, 4.) 350≈6, 5.) 250≈4, and 6.) 150=2.5. In each case the threshold of the cross-correlation estimator was adjusted to keep the probability of a type I error at $1.0X10^{-6}$. Furthermore the resolution of the multitaper estimator was adjusted to maintain the number of tapers at K=7. The following table illustrates the value of the XCR estimator corresponding to the specified probability.

Table 2-Significant correlations by sample size

Number of Samples	Correlation Coefficient (absolute value)
N=550	.222
N=500	.232
N=450	.245
N=350	.277
N=250	.325
N=150	.415

With all 6 sample sizes, the MTM performed better than cross-correlation. The following table illustrates the detection performance by sample size for the two estimators.

Table 3-Correct detections by sample size

Sample Size	Correct Detection (MTM)	Correct Detection (XCR)
550	95%	88%
500	96%	89%
450	92%	84%
350	66%	57%
250	52%	42%
150	7%	3%

The greatest difference between the two methods occurred for the 250-sample data set, where the MTM estimator performed 10 percentage points better than the cross-

correlation estimator. At 150 points (60 seconds=2.5 periods of the stimulus) both methods performed poorly correctly detecting the signal in only 7% (MTM) and 3% (cross-correlation) of the trials respectively.

In all simulations with a white noise background (total of 10), there were 65 false detections with the MTM method. It is important to note that this number represents false detections over all frequencies. There were a total of 3 false detections at the frequencies of interest- .05 cycles/sample and .0167 cycles/sample. In the cross-correlation method there were 19 false detections at the forcing frequencies.

5.2 Application to Actual fMRI Data

While simulations are critical in any study of signal-detection methods, the utility of any data analysis method must be validated by its application to real data. Simulations are well behaved and the conditions under which they are run can be controlled to levels of accuracy not possible in the real world. For example, the assumptions of stationarity and Gaussianity inherent in both the multitaper and cross-correlation methods are typically violated to some degree by actual data. One of the initial problems one encounters when conducting a methodological study of this kind is determining what constitutes a meaningful measure of success or failure. Additionally, with fMRI data, our efforts are hampered by the complexity of the human brain and our elementary understanding of its function. With these considerations in mind, it follows that the best conditions under which to pursue a methodological study of fMRI data is with a simple sensory-processing or motor paradigm. The rationale for this is as follows. Our understanding of elementary sensory and motor processing in the human brain has been helped immeasurably by cortical recordings in both primate and non-primate mammals. With more quintessential

human abilities like memory, emotion, and language, the extrapolation between findings in animals and humans becomes much more tenuous. With a simple motor or sensory-processing paradigm, e.g. finger tapping or watching flashing checkerboards, our a priori notion of where cortical activation should take place is on a much stronger empirical foundation than with "higher" cortical functions like computation and language.

In order to facilitate the comparison between the multitaper and cross-correlation methods with actual data, two experimental paradigms, one visual and one motor, were employed. Both were designed to be simple in design and well localized in their posited foci of neuronal activation. The visual paradigm consisted of flashing checkerboards presented to one visual hemifield. It is known that objects presented to the right of a central fixation point cast images on the nasal retina of the right eye and the temporal retina of the left eye (Sekuler & Blake, 1994). The fibers from both of these retinal areas project to the left primary visual cortex. Therefore, one should expect that a functional acquisition through the primary visual cortex would reveal activation only on the left side. The second paradigm consisted of finger tapping with the right hand only. (thumb to index finger). Movements of fingers, arms, and toes have been shown to be produced by contralateral activation of the primary motor area, M1 (Roland, 1993; Grafton et al., 1991). Furthermore, in a PET study where subjects flexed their finger against a spring, only the contralateral M1 was activated (Roland et al., 1980). The strong lateralization and functional focus of these paradigms make them ideal choices for investigating data-analysis methods.

5.2.1 Visual Experiment

One subject participated in the experiment. Written consent was obtained in accordance with university policies on human subjects. The experiment was conducted using the VA medical center's 3T General Electric (GE) Signa scanner. The standard GE head coil was used for data acquisition. A hi-resolution 3D-volume scan was acquired in order to identify the anatomical area of interest. From that scan, a single 5-mm coronal slice was prescribed 10-mm anterior to the occipital pole. The functional experiment consisted of presenting checkerboards flashing at 3 Hz, alternating in 12-second cycles with a fixation point (a white cross in the middle of the screen). The checkerboards were presented to the right visual hemifield only. The subject was instructed to keep his gaze fixated on the cross at all times and not to turn his gaze toward the flashing checkerboards. Functional data acquisition consisted of 550 128X64 (frequency field-of-view 40 cm; phase field-of-view 20 cm) gradient-echo EPI images with a TR of 400 ms. A 40-degree flip angle was used. The short TR was utilized to combat aliasing effects from respiration and cardiac contributions reported to be a major source of noise in fMRI data (Mitra et al, 1997). A TR of 400 ms allows a frequency of 1.25 Hz to be unaliased. This corresponds to a heartrate of 75 beats/min. It is important to note that the assumption here was that the cardiac frequency is the highest frequency present in fMRI data. This assumption has not been verified experimentally. The subject's heartrate was monitored during data acquisition and the values recorded every 3 seconds. At only two measurements during the acquisition did the subject's heartrate exceed 75 beats/min.

All images were imported into AFNI (Cox, 1995) for analysis. Because anatomical considerations, e.g. the exact locus of activation, were not the primary concern in this

study, functional data was not overlaid onto higher resolution structural scans. The cross-correlation analysis involved 5 reference vectors phase-shifted in the manner described in the simulations. The threshold for both estimators was again held constant at 99.998% (1-1/N) with K=7 for the MTM and $1.0X10^{-6}$ for the cross-correlation estimator. Five reference vectors were used in the cross-correlation analysis. Data processing of the actual data mimicked the final simulation reported above, i.e. 6 data sets with decreasing numbers of samples were constructed by deleting images off the end of the acquisition-550, 500, 450, 350, 250, 150. Results in red will indicate the MTM estimator, while the cross-correlation results will be shown in blue. Figures 4 and 5 show the results for both analysis methods on the 550-image dataset. The MTM estimator found 7 pixels that met or exceeded the .002 level. The majority of these (6/7) were located in the left primary visual cortex (V1). The cross-correlation estimator identified 12 pixels as active, 8 of which were in left V1 In general, there was good agreement between the two methods in terms of their activation maps obtained with 550 samples. The most significant differences were that the cross-correlation method found more extensive activation in the left primary visual cortex and also found some activation in what appears to be the left medial occipital gyrus not found by the MTM estimator. The XCR estimator also identified one pixel active in the right hemisphere. The MTM estimator did not identify any pixels active in the right hemisphere. Six of the 7 pixels identified by the MTM estimator were also identified by the cross-correlation method.

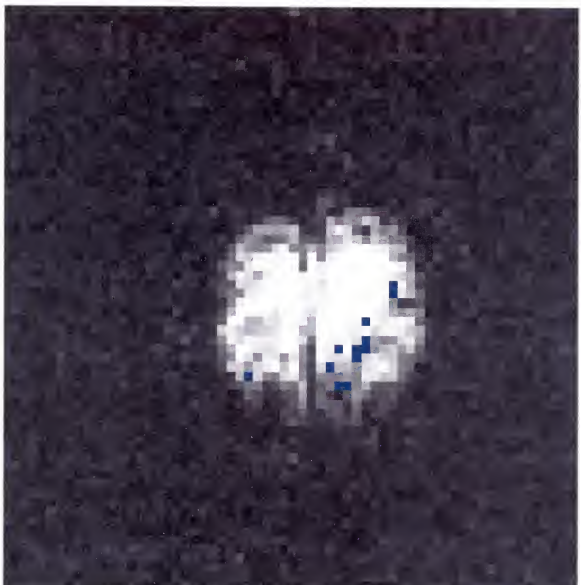

Figure 4- 550-point image xcorr

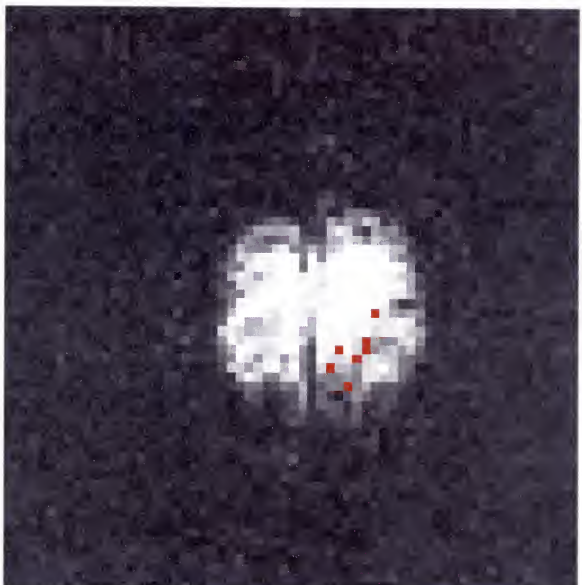

Figure 5- 550-point image MTM

Now consider the effect of truncating the series of images to the first 150. Figures 6 and 7 compare the results for the 150-point series for both estimators. With 150 samples, or 2.5 periods of the stimulus, the MTM estimator found three pixels in the left primary visual cortex active. The cross-correlation estimator found 2 pixels in left V1 active and 3 additional pixels in the brain, one midline, one in the right hemisphere, and an additional pixel in the left hemisphere. Figures 8 and 9 show the pattern of activation found for 250-point images. Both estimators identified 5 pixels as active. The MTM estimator found 4 in left V1, and one in the left hemisphere outside of V1. The cross-correlation estimator found 3 active in V1 and 2 pixels outside of V1, one in the ipsilateral hemisphere and the other in the contralateral. The estimators agreed on 3 of 5 detections. Figures 10 and 11 show the results for the 350-image condition. The MTM estimator identified 5 active pixels in the 350-image condition, all located in left V1. The cross-correlation estimator identified 6 active pixels and 5 of those were in the primary visual cortex. Figures 12, 13, 14, and 15 show the MTM and cross-correlation images for the 450 and 500-image conditions. The trend for the MTM estimator to identify fewer active pixels continued in the 450 and 500-image conditions, with 5 and 6 active pixels respectively compared to 8 and 11 for the cross-correlation estimator. Also, the trend continued for the MTM estimator to have a greater proportion of its pixels in the left V1. Both methods show a linear increase in the number of active pixels as a function of sample size. Furthermore, the percentage of overall detections in the area of the left V1 was high in both methods. The MTM estimator detected fewer pixels overall and had the greatest percentage of pixels in left V1 (93%). The cross-correlation estimator detected 47 pixels overall, 33 of which were in the left V1 (70%).

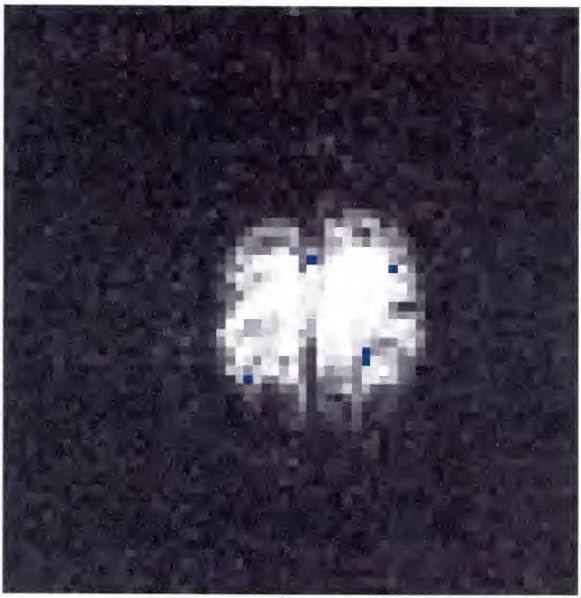

Figure 6- 150-point image xcorr

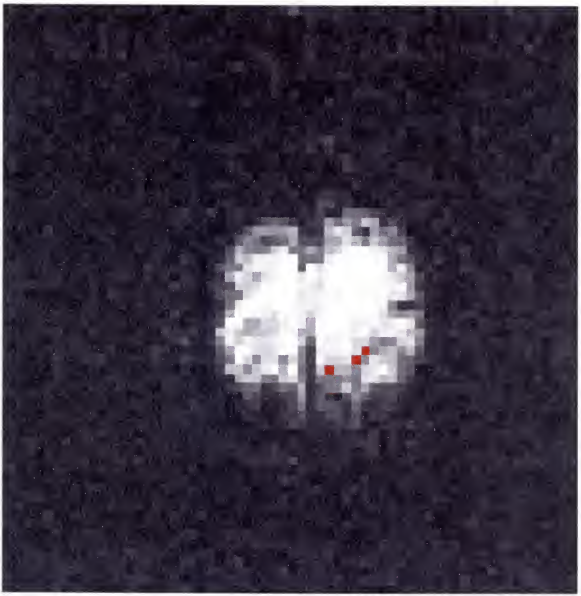

Figure 7-150-point image MTM

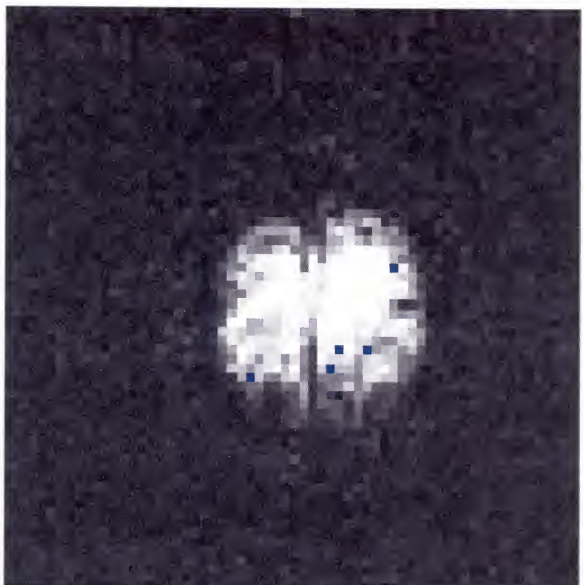

Figure 8-250-point image xcorr

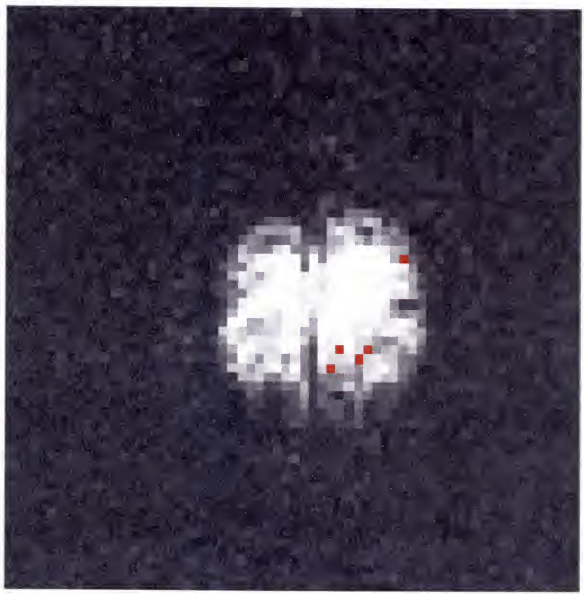

Figure 9-250-point image MTM

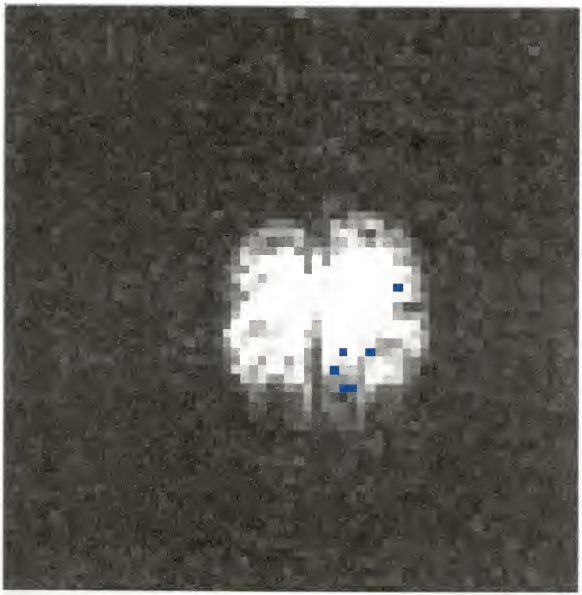

Figure 10-350-point image xcorr

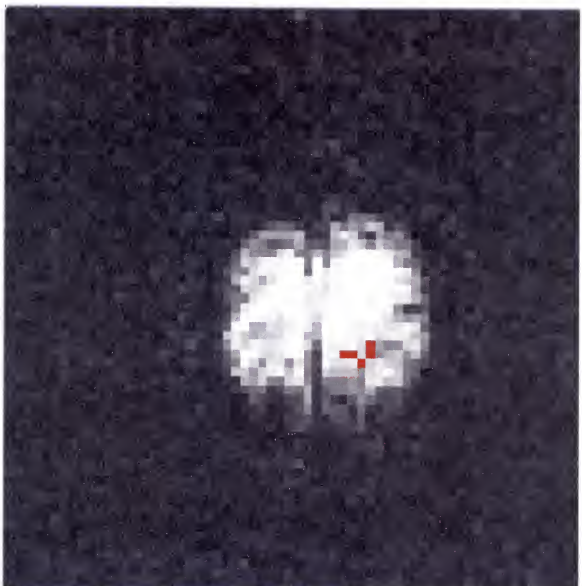

Figure 11-350-point image MTM

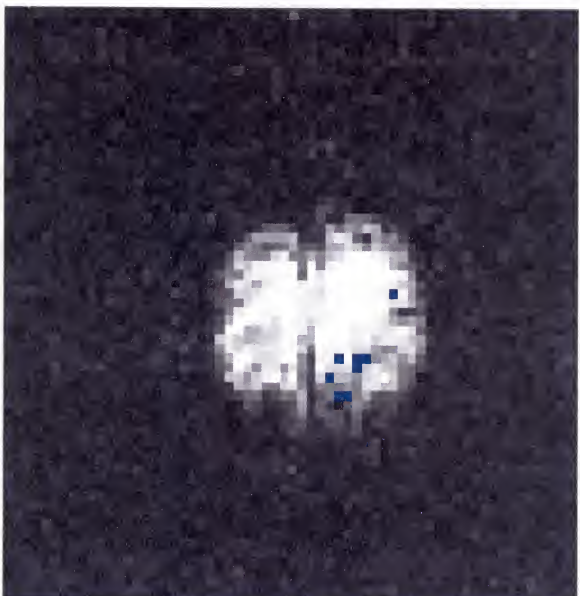

Figure 12-450-point image xcorr

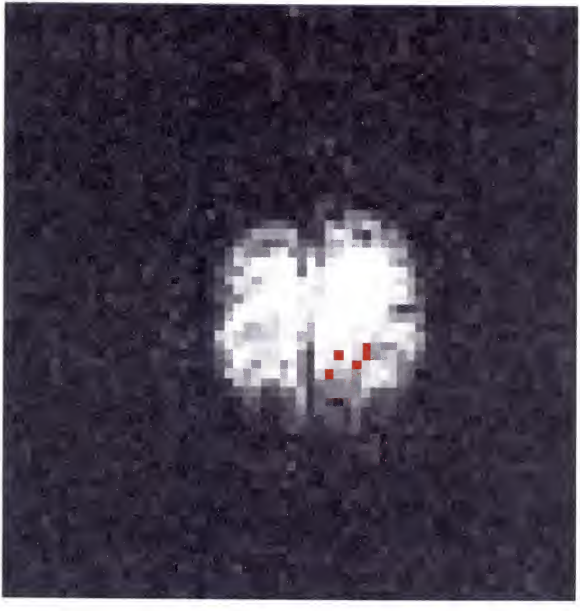

Figure 13-450-point image MTM

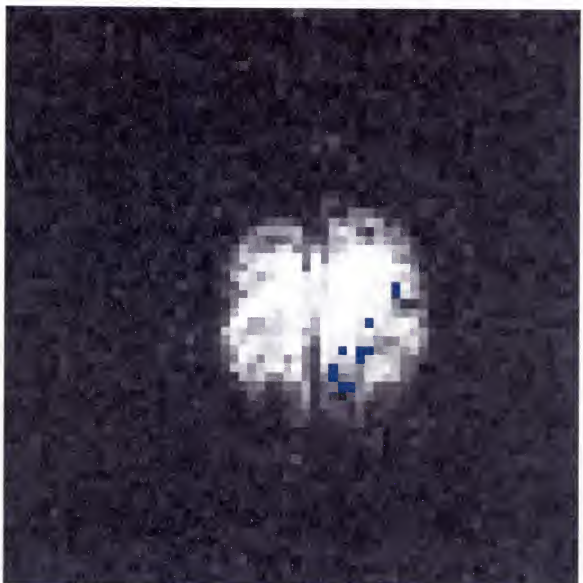

Figure 14-500-point image xcorr

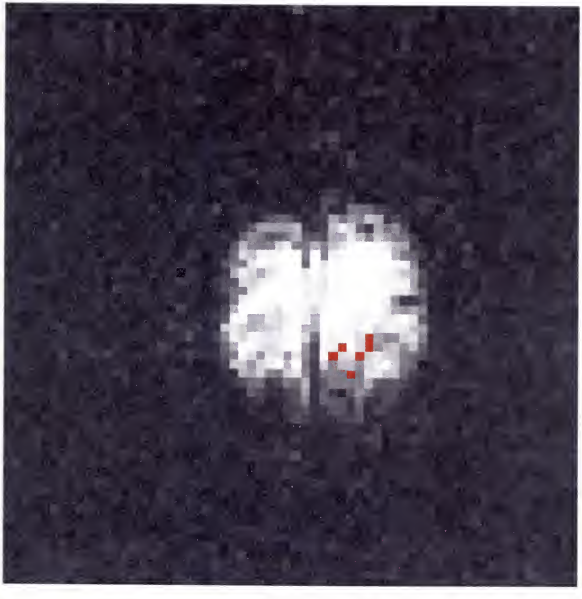

Figure 15-500-point image MTM

The visual experiment was then repeated using the identical experimental paradigm except that the TR was extended to 3 seconds and the flip angle increased to 90 degrees. Seventy-two single-shot EPI images were acquired. The number of stimulus periods was approximately the same as in the 550-image experiment (9 to 9.2). The images for the cross-correlation and MTM are shown in figures 16 and 17.In the slow-sampled visual experiment there were 20 activated pixels identified by the MTM estimator and 33 by the cross-correlation method. Nineteen of the 20 pixels identified by the cross-correlation estimator were also detected by the multitaper method. Seven of the pixels in the MTM image (35%) are clearly outside of left V1 compared to 18 (55%) in the cross-correlation image.

5.2.2 Motor Experiment

One subject participated in the motor experiment. The motor experiment was performed on the same Signa scanner as the visual paradigm. Similar hi-resolution 3D-volume images were used for anatomic localization and prescription. Data acquisition consisted of a single 5-mm 128X64 (frequency FOV 40 cm; phase FOV 20 cm) axial slice prescribed through the area of the motor cortex known to be involved in the control of the hand. As in the visual experiment, single-shot echo-planar imaging was used. A TR of 500 ms was used and a total of 360 images were obtained. This corresponded to 5 periods of the stimulus. A TR of 500 ms allows a frequency of 1 Hz to be well resolved. The heartrate of the subject scanned in this experiment did not exceed 58 beats/minute during the functional acquisition. The subject performed a simple finger-tapping experiment (index finger to thumb) in 36-second periods (18 seconds of tapping followed by 18 seconds of no movement). Figures 18 and 19 show the results for the cross-

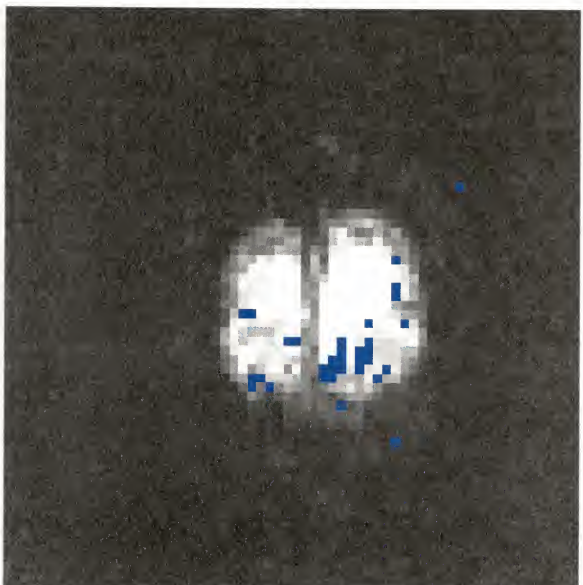

Figure 16-XCR activation map-TR=3

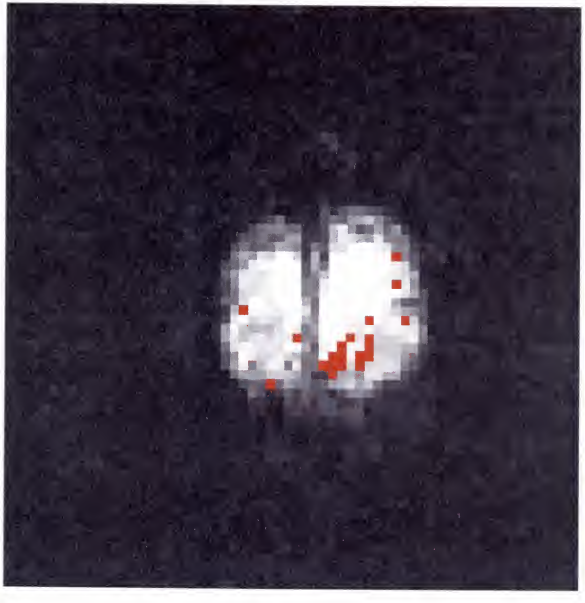

Figure 17-MTM activation-TR=3 seconds

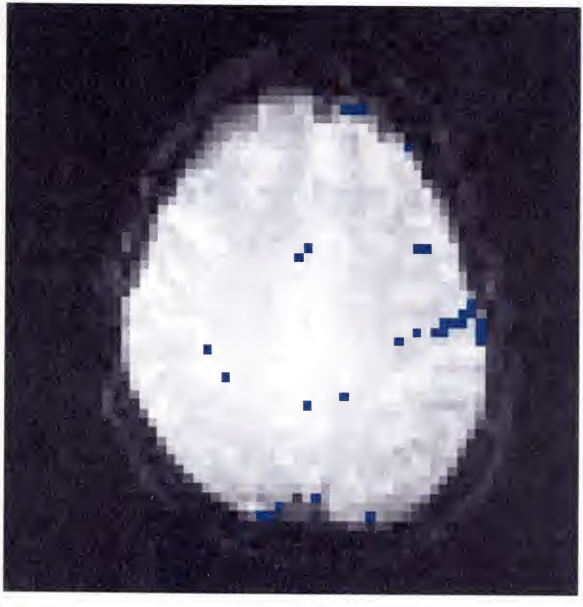

Figure 18- XCR activation right-finger tapping

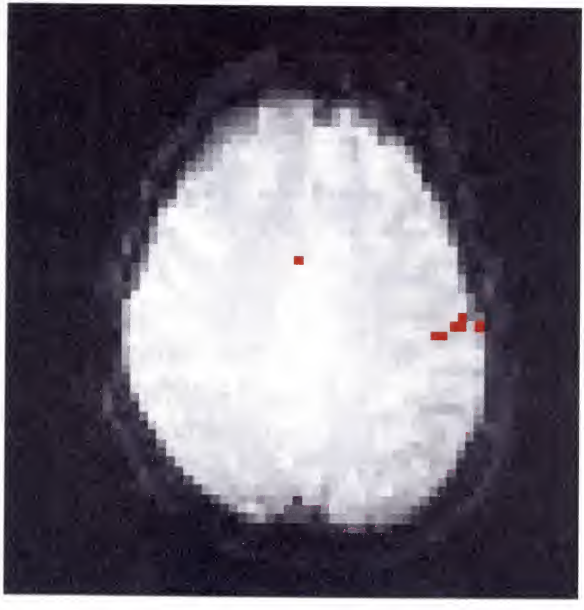

Figure 19- MTM activation right-finger tapping

correlation and multitaper methods. The cross-correlation method found 24 active pixels, 10 of which can be considered to be in the left primary motor area (M1). The multitaper method identified 7 active pixels, 5 of which were in the left M1. All of the pixels identified by the multitaper method were also identified by the cross-correlation estimator.

5.3 Discussion

The purpose of this study was to take an initial step in examining the questions surrounding the efficacy of data-analysis methods in fMRI. Two time-series techniques were compared: a time-domain technique widely employed in fMRI and a frequency-domain technique widely utilized in geophysics, but not currently in fMRI. Comparisons consisted of both simulated and actual fMRI data. Results obtained from the simulations and corroborated by actual fMRI data suggest that the MTM estimator is more robust than the XCR method both in terms of its sensitivity and specificity. In the simulated data, the multitaper method proved a more accurate detector by repeatedly demonstrating more correct detections and fewer false positives. Specifically, in terms of correct detection with fixed sample sizes, the MTM estimator performed 10 to 20 percentage points higher than the cross-correlation estimator. In terms of false-positives, the multitaper estimator also outperformed the cross-correlation estimator. The cross-correlation identified a total of 22 noise-only pixels as active. Compare this with a total of 5 by the MTM estimator. When a series of simulations were performed at a fixed SNR and various sample sizes ranging from 550 to 150, the MTM estimator consistently detected the signal in 4-10% more of the trials. In the first visual experiment where the SNR was held constant and the sample size varied, both estimators showed a linearly

increasing trend in the number of active pixels as a function of sample size. At all sample sizes, the cross-correlation estimator consistently identified more active pixels than the multitaper method. While the majority of the activated pixels found by both methods were in the left primary V1, the XCR method produced a greater percentage of pixels (30%) that were clearly outside of the left V1. Only 7% of the pixels identified by the MTM estimator were outside of the left V1. The motor study revealed a similar pattern, with the MTM estimator detecting fewer overall active pixels (7 vs. 24). However, 5 of the 7 pixels (71%) identified by the multitaper method were in the left M1 as compared to 10 of the 24 (41%) with the XCR method. In other words, more than half of the pixels identified as active by the XCR estimator in a simple finger-tapping experiment were outside of primary motor cortex. The pattern was also repeated in the slow-sampled visual experiment where the MTM estimator demonstrated a greater percentage of active pixels in the primary sensory area. In fact, in every actual dataset, the XCR method produced more apparent false positives than the MTM estimator. Considered together, the results from the simulations and the actual fMRI data suggest that the pixels identified by the MTM estimator represent a more accurate picture of periodic fluctuations at the forcing frequency than those found by the XCR method. The natural extrapolation of this finding is that the MTM estimator provides a more reliable picture of the underlying neuronal activity. To reiterate the support for this contention, recall that the multitaper method proved itself a more reliable detector in 11 simulations with various numbers of sample points and SNRs. The majority of the sample sizes employed and all of the SNRs are indicative of those found in actual fMRI data. Additionally, in the actual fMRI data, the multitaper method consistently showed activation patterns that were more restricted

to the a priori posited areas of neuronal activation. The XCR method similarly showed activation in the appropriate neuronal areas, but it also indicated a larger percentage of pixels (from 30% to 60%) outside of the primary sensory areas. Contrast this with the percentages of false positives found with MTM (7% to 35%).

There are several plausible reasons for the superior performance of the multitaper method. First, the multitaper method is a frequency-domain approach. As previously mentioned, statistical estimation in the frequency domain inherently holds several advantages over time-domain estimation. Although the Fourier transform of a signal represents a trivially sufficient statistic (no information is lost), adjacent values (frequencies) in the Fourier transform can often be regarded as uncorrelated. This statement is true by definition if the process is stationary. On the other hand, adjacent values in the time domain usually cannot be regarded as uncorrelated. If random variables (RV) can be shown to be uncorrelated, then statistical inference is greatly facilitated. The reason for this is that joint probabilities, as well as the means and variances of linear combinations of those RVs (as one encounters with additive noise) can be more easily calculated. Secondly, since the Fourier transform in the discrete-time case is the summation of many terms, the resulting distribution can often be considered Gaussian by a weak central limit theorem argument regardless of the distribution of the original data sequence. Further, the use of the orthogonal (uncorrelated) discrete prolate spheroidal sequences to window the data allows for the degrees of freedom in the spectral estimate to be increased considerably. This produces a substantial narrowing of confidence intervals for point estimation. Finally, harmonic analysis (the detection of sinusoids) in the MTM is accomplished via the theory of complex least squares and multiple

regression. Least squares and multiple regression are invaluable statistical estimation procedures supported by a body of research that dates back more than a century (Draper & Smith, 1998). No such body of work supports the XCR estimator. The use of multiple regression and the discrete-time/discrete-frequency Fourier transform allow for the precise specification of the probability of at least one false detection over all possible Fourier frequencies. In the cross-correlation method, the use of multiple reference vectors, which are clearly not uncorrelated, and the lack of knowledge of the underlying sampling distribution make accurate specification of the probability of a type I error prohibitively difficult. This is especially true of the multiple-comparison problem, since the reference vectors are not uncorrelated. Without solid probabilistic and inferential ground for setting the threshold in the XCR method, the danger is that the investigator simply adjusts the threshold until the data fit the researcher's preconceived notions of how the activation maps should appear. Also, the MTM estimator naturally provides information on all the frequencies present in the sample from 0 to the Nyquist. With properly sampled data, this has the potential to provide valuable information on a wide range of oscillatory phenomena that may occur as the result of physiologic processes or data acquisition. This can potentially yield valuable information for the researcher when choosing the frequency of stimulation. For example, there could very well be periodic phenomena that occur in the functional data as a result of the use of a particular pulse sequence on a particular machine, than the MTM estimator could identify these phenomena and allow the researcher to avoid that frequency in their experimental paradigm. Note that this is done routinely in EEG experiments where 60 Hz interference (in the U.S.) due to polarity-switching electrical currents is commonplace. In those

studies, stimulation frequencies that are multiples of 60 Hz are routinely avoided because of the possibility of contamination.

Finally, this study also raises awareness of a serious issue normally not discussed in the fMRI literature- the issue of detector performance on the reliability of fMRI as a clinical and experimental tool. This issue must be addressed if fMRI is to advance as a human brain-mapping technique. Normally discussions on repeatability/reliability in fMRI center around pulse sequences, scanner stability, subject movement, physiologic noise, the coupling between neuronal activity and the cerebral blood flow, and neuronal habituation. This study suggests that the statistical stability of the estimator may in fact be one of the largest sources of variation. The fact that the multitaper method and the XCR estimator proceed from identical assumptions about the underlying signal and yet can generate very different activation maps when analyzing the same set of data underscores the need to bring this issue to the forefront in fMRI research. The issue of detector performance has been central to research in areas as seemingly dissimilar as electrical engineering and psychology for decades. Given the large numbers of psychologists conducting human brain-mapping research, it is puzzling why its discussion in fMRI has been so limited to date.

In terms of limitations, it is probable that the correct detection rates and the false positive rates found in the simulations are optimistic for both estimators. For example, the large number of simulations with white Gaussian noise generated relatively few false positives for both estimators given that the matrix sizes were 64X64. Note that this means that there were 4096 pixels per image and over 50 simulations that represents 204,800 pixels. Experience with actual fMRI data would suggest that the false positive rate should

be higher. Since we are currently limited by our knowledge of the true nature of the noise characteristics in fMRI, it is likely that the white and red noise used in the simulations are simplistic approximations to the noise environment in actual fMRI data. Also, it is likely that the presence of additional line components in actual fMRI data is almost a certainty given the literature and this could easily complicate the analysis for both methods, i.e. increase the probability of false positives. Still, while the rates may be optimistic for both estimators, they were substantially lower for the MTM estimator as compared to the XCR method. In addition to our incomplete understanding of the signal in fMRI, which leads to imprecise modeling, there are other limitations in this study. Firstly, for time/frequency based analysis in general, the slow nature of the hemodynamic response presents a problem. It has been shown that the BOLD effect in periodic stimulation paradigms begins to drop off dramatically as the forcing frequency falls below a period of 12 seconds (Bandettini et al, 1999). Recall that the delay of the response is on the order of a few seconds. It is intuitively clear that if the period of stimulation were on the order of 12 seconds (6 on, 6 off), the cerebral blood flow response would not have sufficient time to increase and decrease in response to the stimulus. As a result, experimental paradigms typically utilize forcing frequencies on the order of 1/60, 1/30, or 1/24. The problem for a time/frequency analysis of these frequencies lies in their proximity to zero. An important theorem in statistics know as the Cramer-Rao inequality can be used to show that the minimum variance of frequency estimators increases dramatically as one approaches zero or the Nyquist frequency (Kay, 1998). Furthermore, if the AR(1) characterization of fMRI noise is correct, then the experimental frequencies would typically fall in the part of the spectrum with the most

energy. It would be beneficial if the frequency of interest could be separated somewhat from the high-energy low frequencies for detection. It is unclear how this problem can be ameliorated. The slow-cycling nature of the BOLD effect makes frequencies higher than 1/12 Hz difficult to obtain. Another possibility would be to modulate the signal to another part of the frequency spectrum where signal detection would be more favorable from a minimum-variance viewpoint. However, the noise in that part of the spectrum will also be modulated along with the signal, so it is unclear how much advantage would be gained. Finally, both data analysis methods employed here, as well as those mentioned earlier (t-test, Kolmogorov-Smirnov test), are univariate procedures, where each pixel is treated as its own independent time series. Clearly, this is a simplification for the purposes of analysis. The data obtained in fMRI is an example of a multivariate time series. Multivariate time-series analysis techniques are disproportionately more difficult than univariate series for many reasons (Thomson & Chave, 1991). It should be noted that the MTM technique can be extended to the multivariate case (Thomson & Chave, 1991). These and other multivariate techniques should be investigated for their utility in fMRI data analysis. Furthermore, the variant of the multitaper method employed here is only one of many techniques available for the analysis of both univariate and multivariate time-series data. It is likely that MTM techniques specifically designed for non-stationary and non-linear time-series analysis would be better suited to actual fMRI data (Thomson, 1999).

The results presented in this study support that the MTM estimator is more robust than the cross-correlation estimator in correctly detecting periodic signal changes in fMRI data. Furthermore, these results also suggest that the multitaper method is more resistant

to false detections. While intriguing, these data are preliminary and must be expanded in several dimensions before any definitive results can be stated. The MTM method should also be compared with other data analysis methods currently employed in time-series analysis to find out which methods are the most effective detectors. It may turn out that the most accurate and reliable way to analyze fMRI data is to use more than one detection method on each dataset and accept only those pixels, which are found by more than one estimator. While this is undesirable from a computational and time perspective, it may necessary to ensure accurate activation maps for clinical applications, i.e. functional mapping prior to neurosurgical intervention, where reliable data is a matter of patient safety.

5.4 Future Directions

As a continuation of this study, future work is planned in several areas. First, data will be collected from both phantom and biological sources under a variety of fast-sampling conditions in order to more closely exam the nature of fMRI noise (in the absence of any exogenous stimulation). This knowledge is crucial for experimental design, data acquisition, and analysis. Secondly, different formulations of the MTM technique including multivariate models, jacknifed MTM estimates, and multitaper techniques for non-stationary processes (Thomson, 1999) will be investigated to determine their applicability to fMRI data analysis. Thirdly, work is planned to examine the effects of altering the forcing frequency both by modulating it to a more favorable portion of the spectrum and by attempting to collect the BOLD response at higher experimental frequencies. Finally, further work is planned to compare the MTM estimator with other

univariate and multivariate procedures to gain a deeper understanding on how the statistical performance of the estimator affects the activation maps obtained in fMRI.

REFERENCES

Astrup, J., Moller-Sorenson, P., & Rahbek-Sorenson, H. (1981). Inhibition of cerebral oxygen and glucose consumption in the dog by hypothermia, pentobarbital, and lidocaine. *Anesthesiology,* 55, 263-268.

Atkins, P.W. (1986). *Physical chemistry.* 3rd ed. W.H. Freeman:New York.

Bandettini, P.A., Kwong, K.K., Davis, T.L., Tootell, R.B.H., Wong, E.C., Fox, P.T., Belliveau, J.W., Weisskopf, R.M., & Rosen, B.R. (1997). Characterization of cerebral blood oxygenation and flow changes during prolonged brain activation. *Human Brain Mapping,* 5,93-109.

Bandettini, P.A., Wong, E.C., Hinks, R.S., Tikofsky, R.S., Hyde, J.S. (1993). Time course EPI of human brain function during task activation. *Magnetic Resonance in Medicine,* 25, 390-397.

Baron, J.C., Rougemont, D., Soussaline, F., Bustany, P., Crouzel, C., Bousser, M.G. & Comar, D. (1984). Local interrelationships of cerebral oxygen consumption and glucose utilization in normal subjects and in ischemic stroke patients: a positron tomography study. *J. Cereb. Blood Flow Metab.* 4,140-149.

Borowsky, W.I., & Collins, R.C. (1989). Metabolic anatomy of brain: a comparison of regional capillary density, glucose metabolism, and enzyme activities. *J. Comp. Neurol.* 288, 401-413.

Cohen, J.D., Forman, S.D., Braver, T.S., Casey, B.J., Servan-Schreiber, D., & Noll, D.C. (1994). Activation of the prefrontal cortex in a nonspatial working memory task with functional MRI. *Human Brain Mapping,*1,293-304.

Cotman, C.W., Monaghan, D.T., Otterson, O.P., & Storm-Mathisen, J. (1987). Anatomical organization of excitatory amino acid receptors and their pathways. *Trends in Neuroscience,* 10(7), 273-280.

Cox R.W., Jesmanowicz, A., & Hyde J.S. (1995). Real-time functional magnetic resonance imaging. *Magnetic Resonance in Medicine,* 33, 230-236.

Denison, D.G.T., Walden, A.T., Balogh, A., & Forsyth, R.J. (1998) Multitaper testing of spectral lines and the detection of the solar rotation frequency and its harmonics. *Technical Report TR-98-04,* Statistics Section, Imperial College, London.

Demaerel, P., Heiner, L., Robberecht, W., Sciot, R., & Wilms, G. (1999). Diffusion-weighted MRI in sporadic Creutzfeldt-Jakob disease. *Neurology*, v52, p205-8.

DeYoe, E.A., Bandettini, P., Neitz, J., Miler, D. & Winans, P. (1994). Functional magnetic resonance imaging of the human brain. *Journal of Neuroscience Methods*, 54,171-187.

Dunning, H.S., & Wolff, H.G. (1937). The relative vascularity of various parts of the central and peripheral nervous system of the cat and its relation to function. *J. Comp. Neurol.*, 67, 433-450.

Freund J.E. (1992). *Mathematical statistics.* 5th edition. Prentice Hall, NJ.

Fullerton, G. D. (1988). Physiologic basis of magnetic relaxation. In D.A. Stark & W.G. Bradley, JR. (Eds.) *Magnetic resonance imaging.* Mosby:St. Louis.

Furchgott, R.F., & Zawadzki, J.V. (1980). The obligatory role of endothelial cells in the relaxation of arterial smooth muscle by acetylcholine. *Nature*. 288,373-376.

Hand, P.J., Greenberg, J.H., Sylvestro, A., Weller, L., & Reivich, M. (1979). A normal and developmentally local glucose utilization with natural stimlation of a single receptor organ. *Act. Neurol. Scand. Supl.*, 72, 60, 46-477.

Haykin, S. & Thomson, D.J. (1999). Signal detection in a nonstationary evironment reformulated as an adaptive pattern classification problem. *Proceedings of the IEEE*, 86,2325-2344.

Kuschinsky, W., Klein, B., Schroeck, H., & Vetterlein, F. (1987). The local density of perfused capillaries in the rat brain. *J. Cereb. Blood Flow. Metab.* 7(supp1):S333.

Kwong, K.K., Beliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy D.N., Hoppel, B.E., Cohen, M.S., Turner, R., Cheng, H.M., Brady, T.J., Rosen, B.R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc.Natl.Acad.Sci.USA*, 5675-5679.

Larrabee, M.G. (1958). Oxygen consumption of excised sympathetic ganglia at rest and in activity. *Journal of neurochemistry*, 2:81-101.

Mann, M. & Park, J. (1994). Global-scale modes of surface temperature variability on interannual to century timescales. *American Geophysical Union*, 1-29.

Mata, M., Fink, D.J., Gainer, H., Smith, C.B., Davidsen, L., Savaki, H., Schwartz, W.J., & Sokoloff, L. (1980). Activity-dependent energy metabolism in rat posterior pituitary primarily reflects sodium pump activity. *J. Neurochem.* 34(1), 213-215.

McIlwain, H. (1954). Study of human cerebral biopsy specimens in an electrically excited condition. *Arch.Neurol.Psychiatric.*, 71, 488-495.

Merzbacher, E. (1961). *Quantum mechanics.* Wiley:New York

Miller, D.H., Grossman, R.I., Reingold, S.C., McFarland, H.F. (1998). The role of magnetic resonance techniques in understanding and managing multiple sclerosis. *Brain,* v121, p3-24.

Miller, K.S. (1974). *Complex stochastic processes.* Addison-Wesley, Mass.

Mitra, P.P. & Pesaran, B. (1999). Analysis of dynamic brain imaging data. *Biophysical Journal*, In Press.

Mitra, P.P., Ogawa, S., Hu X., & Ugurbil, K. (1997) The nature of spatiotemporal changes in cerebral hemodynamics in functional magnetic resonance imaging. *Magnetic Resonance in Medicine,* 37, 511-518.

Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proc.Natl.Acad.Sci.USA*, 89, 5951-5955.

Park, J., Lindberg, C.R., & Vernon, F.L. (1987). Multitaper spectral analysis of high-frequency seismograms. *Journal of Geophysical Research,* 92, 12675-12684.

Percival, D.B. & Walden, A.T. (1993) *Spectral analysis for physical applications.* Cambridge Univeristy Press.

Priestly, M.B. (1981). *The spectral analysis of time series.* John Wiley & Sons.

Purdon P.L. & Weisskopf, R.M. (1998). Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Human Brain Mapping,* 6,239-249.

Roland, P. (1993). *Brain activation.* John Wiley & Sons.

Slepian, D. (1976) On bandwidth. *Proceedings of the IEEE,* 64, 292-300.

Slepian, D. (1978). Prolate spheroidal functions, Fourier analysis, and uncertainty-V:The discrete case. *The Bell System Technical Journal,* 57, 1371-1430.

Slepian, D. & Pollak, H.O. (1961). Prolate spheroidal wave functions, Fourier analysis and uncertainty-I. *Bell System Technical Journal,* 40, 43-63.

Slichter, P. (1978). *Principles of magnetic resonance.* Wiley:New York.

Stoica, P. & Moses, R. (1998). *An introduction to spectral analysis.* Prentice Hall:NJ.

Thomson, D.J. (1994). An overview of multiple-window and quadratic-inverse spectrum estimation methods. *IEEE international conference on speech, acoustics, and signal processing,* 185-194.

Thomson, D.J. (1990a). Quadratic-inverse spectrum estimates: applications to palaeoclimatology. *Phil. Trans. R. Soc. Lond.,* 332,539-597.

Thomson, D.J. (1990b). Time series analysis of Holocene climate data. *Phil.Trans.R.Soc.Lond.,* 601-616.

Thomson, D.J. (1982). Spectrum estimation and harmonic analysis. *Proceedings of the IEEE,* 70, 1055-1094.

Thomson, D.J. & Chave, A.D. (1991). Jacknifed error estimates for spectra, coherences, and transfer functions. In Simon Haykin, ed. *Advances in Spectrum Analysis and Array Processing,* vol I, 58-109.

Tyson, G., Gross, P., Tajima, T., Nakata, H., Pettigrew, K., Patlak, C., & Fenstermacher, J. (1987). Local cerebral blood flow, size of capillary beds, and velocity of capillary flow in seven brain structures. *J. Cereb. Blood Flow Metab.* 7(supp1):S341.

Wehrli, F. (1988). Principles of magnetic resonance. In D.A. Stark & W.G. Bradley, JR. (Eds.) *Magnetic resonance imaging.* Mosby:St. Louis.

Zheng, D., LaMantia, A.S., & Purves, D. (1991). Specialized vascularization of the primate visual cortex. *J. Neurosci.* 11(8), 2622-2629.

BIOGRAPHICAL SKETCH

Wayne M. King received his bachelor's degree from Kent State University in 1983. He then served 6 years on active duty in the United States army before returning to graduate school at the University of Florida in 1992. He received his Master's degree from the University of Florida in 1994 and his Ph.D. in audiology from the University of Florida in 1999.

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Carl C. Crandell, Chairman
Associate Professor of
Communication Sciences and
Disorders

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Stephen J. Blackband
Associate Professor of Neuroscience

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Kenneth Gerhardt
Professor of
Communication Sciences and
Disorders

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Christiana M. Leonard
Professor of Neuroscience

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Linda J. Lombardino
Professor of
Communication Sciences and
Disorders

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Murali Rao
Professor of Mathematics

This dissertation was submitted to the Graduate Faculty of the Department of Communication Sciences and Disorders in the College of Liberal Arts and Sciences and to the Graduate School and was accepted as partial fulfillment of the requirements for the degree of Doctor of Philosopy.

August 1999

Dean, Graduate School